The Lattice Boltzmann Method for Wind Farm Simulations: Validation and Application

Henry Korb

UPPSALA
UNIVERSITET

Dissertation presented at Uppsala University to be publicly examined in E45, Cramérgatan 3, Visby, Monday, 25 September 2023 at 10:00 for the degree of Licentiate of Philosophy. The examination will be conducted in English. Chairman of the grading committee: Associate Professor Martin Hansen (Technical University of Denmark).

Abstract
Korb, H. 2023. The Lattice Boltzmann Method for Wind Farm Simulations: Validation and Application. 60 pp. Uppsala: Uppsala University.

Many new challenges in wind energy require the use of large eddy simulation for accurate modeling of wind farm flows. However, the immense computational cost hinders its use in research and industry. The lattice Boltzmann method is the most promising candidate to date to achieve the highest level of accuracy while decreasing computational cost by orders of magnitude compared to traditional methods. In this thesis, I present further development of the lattice Boltzmann method for wind energy and compile various applications, such as industrial use, generation of training data for machine learning, and analysis of wind farm control paradigms. In order to evaluate the requirements of di#erent industrial stakeholders, we conduct a survey among industry experts on the use of large eddy simulation and show that the run time requirements indicated by many respondents can be met with the current state of the lattice Boltzmann method. In a validation study, the lattice Boltzmann method is as accurate as traditional Navier-Stokes solvers, while reducing computational cost by one to more orders of magnitude. A convolutional neural network is trained to predict average flow velocities in the wake of a single turbine. The predictions exhibit very high accuracy at execution times similar to engineering models. The lattice Boltzmann method enables the generation of larger training sets at a feasible computational cost. A proof of concept is provided for the use of reinforcement learning to discover new, cooperative wind farm control mechanisms. In an extensive analysis of the helix approach, its physical mechanisms are elucidated and a thorough parameter study of the wake of a single turbine is provided. We also study the interaction of multiple helical wakes, providing a way to extend the approach from a pair of turbines to wind farms.

Henry Korb, Department of Earth Sciences, Wind Energy, Villav. 16, Uppsala University, SE-75236 Uppsala, Sweden.

© Henry Korb 2023

URN urn:nbn:se:uu:diva-510666 (http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-510666)

List of papers

This thesis is based on the following papers, which are referred to in the text by their Roman numerals.

I Asmuth, H., Korb, H., Ivanell, S., 2023. How Fast is Fast Enough? Industry Perspectives on the Use of Large-eddy Simulation in Wind Energy. J. Phys.: Conf. Ser. 2505, 012001.

II Korb, H., Asmuth, H., Ivanell, S., 2023. Validation of a Lattice Boltzmann Solver Against Wind Turbine Response and Wake Measurements. J. Phys.: Conf. Ser. 2505, 012008.

III Asmuth, H., Korb, H., 2022. WakeNet 0.1 - A Simple Three-dimensional Wake Model Based on Convolutional Neural Networks. J. Phys.: Conf. Ser. 2265, 022066.

IV Korb, H., Asmuth, H., Stender, M., Ivanell, S., 2021. Exploring the application of reinforcement learning to wind farm control. J. Phys.: Conf. Ser. 1934, 012022.

V Korb, H., Asmuth, H., Ivanell, S., 2023. The characteristics of helically deflected wind turbine wakes. Journal of Fluid Mechanics 965, A2.

Reprints were made with permission from the publishers.

Division of the work among authors

Paper I Henrik Asmuth (HA) developed the survey with input from Henry Korb (HK) and Stefan Ivanell (SI). The survey results were analyzed by HA. The adaption of the simulation software was done by HK and HA. The manuscript was written by HA with input by HK and SI.

Paper II HK conducted the simulation and analyzed the results. HA acquired the data from the benchmark, HA and HK acquired permissions and data from other benchmark participants. HK wrote the manuscript with input by HA and SI.

Paper III HK and HA developed and implemented the model, HK organized the publication of code and training data while HA wrote the manuscript with input by HK.

Paper IV HK, HA and Merten Stender (MS) developed the study. HK implemented the reinforcement learning environment in `elbe` under guidance from HA and with input from MS. HK wrote the manuscript with input from HA, MS and SI.

Paper V The theoretical analysis was done by HK. The numerical study was designed by HK with input from HA and SI. HK conducted and analyzed the simulations with input from HA. HK wrote the manuscript with input from HA and SI.

Contents

1 Introduction 7
- 1.1 Wind Farm modeling 8
- 1.2 Wind Farm Control 9
- 1.3 Machine Learning 9
- 1.4 Aim of this work 10
- 1.5 Outline 10

2 Methodology 12
- 2.1 The Lattice Boltzmann Method 12
  - 2.1.1 Boundary Conditions 15
  - 2.1.2 Large Eddy Simulation 16
  - 2.1.3 The Actuator Line Method 17
- 2.2 Artificial Neural Networks 18
- 2.3 Reinforcement Learning 20

3 Summary 22
- 3.1 Industry perspectives on the application of LES 22
- 3.2 Validation against full scale measurements 26
- 3.3 WakeNet 0.1 - A convolutional neural network for wind turbine simulations 31
- 3.4 Discovering New Control Strategies with Reinforcement Learning 34
- 3.5 Characterizing the Helix Approach 40

4 Conclusion and Outlook 48

5 Samanfattning 51

1. Introduction

The wind energy sector remains the fastest-growing provider of renewable energy as documented by IEA, 2022. However, according to Veers et al., 2019, accurately predicting the power production, loads, and wake effects of wind turbines throughout the turbines' lifetime still remains one of the great challenges of research in this field. Wind turbines are situated in the atmospheric boundary layer (ABL) and thus exposed to changing weather conditions as well as turbulence. Due to clustering in wind farms, turbines are also subjected to wake effects of upstream turbines, leading to drastic reductions in power production and enhanced degradation due to higher turbulence intensity, discussed for example in Lee et al., 2013.

Accurately modeling the aforementioned effects requires high-fidelity methods of the flow in and around wind farms. High-fidelity methods solve approximations of the governing equations and include fewer simplifications than lower fidelity methods. Consequently, higher fidelity methods also provide more detailed information. A variety of different methods have to be combined to simulate the flow in wind farms. First, an accurate method to simulate the ABL is needed, including appropriate models for the bulk flow, the surface layer, and turbulence. Additional models are needed if other features of the environment, such as forests and complex terrain, are to be considered. The turbine, including loads and the controller, have to be modeled as well. Due to the advent of large wind farms, the interest in wind farm control has increased significantly over the last few years, posing new challenges in wind farm modeling. New advancements in wind turbine technology also drive the need for new modeling capabilities, such as the development of floating offshore wind turbines and the trend towards longer, more flexible blades.

The method with the highest computational fidelity feasible today is large eddy simulation (LES). While being well established as a tool in research, the industrial application of LES is limited, mainly due to the immense computational cost of conducting a single simulation. In recent years, a new method for conducting LES has been introduced to the field of wind energy research, the lattice Boltzmann method (LBM). The LBM has the potential to reduce computational cost to a level that LES becomes feasible to be deployed in the industrial sector and to be used in new applications, which are prohibitively expensive with conventional methods. The reduction in computational cost is enabled by the use

of graphics processing units (GPUs), which can be used very efficiently with LBM.

1.1 Wind Farm modeling

The application of LES to simulations of the ABL dates back to Deardorff, 1972 and has been used for a large number of fundamental studies, such as Shaw and Schumann, 1992 or Khanna and Brasseur, 1997. In their review of LES of the ABL, Stoll et al., 2020 find a growing number of publications using LES to study the ABL, applied to virtually all areas of research. They attribute the trend towards the use of LES to the development of modern computing, which enables the use of larger numbers of grid points to increase resolution or the size of the studied domain. The study also highlights that LES has become a trusted tool due to a focus on theory, model development, and thorough validation efforts.

The use of LES to examine the flow around wind turbines was enabled by the development of the actuator disk and line methods, introduced in Sørensen and Myken, 1992 and Sørensen and Shen, 1999, respectively. Actuator methods model the turbine as body forces based on blade-element/ momentum theory (BEM). While it is also possible to simulate rotors as solid objects, referred to as full-rotor simulations, the grid size required to resolve the boundary layer near the blades usually results in prohibitively high computational costs. R. F. Mikkelsen, 2003 conducted fundamental examinations of both methods, while Troldborg, 2009 focused on the actuator line method. Fundamental investigations, for example into tip spirals and simulating while wind farms were conducted by Ivanell, 2009. Sanderse et al., 2011 give an early review of computational fluid dynamics of wind farms and conclude that the actuator line is the more accurate actuator method, while the actuator disk is in wider use due to lower computational cost. Both methods are deemed appropriate to model the far-wake region. The authors also emphasize the improved accuracy over lower fidelity models. Mehta et al., 2014 note a lack of cross-comparison and benchmarking with common experimental data and the absence of error quantification of LES. The survey of modeling methods by Breton et al., 2017 finds that LES can provide reliable results for quantities of interest in wind farm flows. Porté-Agel et al., 2020 give an overview of wind farm flows and also attest to the accuracy of LES. They also discuss the effect of wind turbine wake on lifetime power production of wind turbines.

1.2 Wind Farm Control

The field of wind farm control seeks to maximize power production as well as turbine lifetime by mitigating the effects of wind turbine wakes on a farm level. Recently, other objectives linked to wind farms' changing role in the energy system have also attracted interest from researchers, for example, Singh et al., 2022. The approaches are usually categorized into two groups, axial induction control and wake steering. The review by Kheirabadi and Nagamune, 2019 found that axial induction control did not conclusively provide power improvement but effectively improved turbine lifetime, while wake steering showed consistent increases in power production at the cost of reduced turbine lifetime. The authors also stress the need for high-fidelity simulations to accurately assess the effects of wind farm control, as lower fidelity models often reported improvements which then could not be reproduced by higher fidelity models. In recent years, dynamic approaches have garnered attention, their theoretical potential proved first by Goit, Munters, et al., 2016. A more applicable approach, named the *pulse*, was presented by Munters and Meyers, 2018, consisting of a sinusoidal variation of the collective pitch, effectively modulating the induction of the turbine. The approach was developed to mimic the behavior of the optimal controller developed in Munters and Meyers, 2017. The pulse was followed by the *helix* approach, presented by Frederik et al., 2020b, which creates a helical wake through individual pitch control. The authors show that the helical wake leads to a strong increase in wake recovery, resulting in increases in power production of over 40 % by a downstream turbine. The method has attracted much attention due to its simplicity and effectiveness. Frederik et al., 2020a compared different dynamic control approaches and found that the helix approach is a good compromise between increasing wake mixing and increased variation of the turbine loading. Muscari et al., 2022 conducted a dynamic mode decomposition of a helical wake and Berg et al., 2022 applied the helix approach to floating wind turbines. Taschner et al., 2023 recently examined the effect of the amplitude of the helix on power and damage equivalent loads.

1.3 Machine Learning

Data-driven methods in general and machine learning methods specifically are now an important tool in every field of research, and the wind energy field is no exception. Machine learning has been applied to, among others, wind farm modeling and wind farm control. Many machine learning methods rely on artificial neural networks. A neural network consists of connected layers of nodes that perform some transformation on its input. Different types of layers exist, for exam-

ple convolutional layers operating on pixelated data or recurrent layers, that retain information throughout multiple invocations of the network. Feed-forward neural networks take some input data, apply an activation function on the weighted input, and return the result. Neural networks have to learn a set of parameters, for example the weights in a feed-forward network, through training. In supervised learning, the network is trained on data where the correct value is known, while reinforcement learning seeks to optimize a cumulative reward.

Many varieties have been tried to accomplish wind farm modeling. Most often, specific quantities of interest are predicted, such as turbine power output, see for example Sun et al., 2020. Others, such as Li et al., 2022, train a convolutional neural network on LES data to predict the evolution of the flow field in a wind farm over time. Both report highly accurate results, however, generating or acquiring training data presents a major limiting factor and the networks are typically not reusable for different wind farms.

Machine learning models are especially interesting for wind farm control applications as they allow rapid generation of states of the wind farm to be used in optimization algorithms for wind farm control. Machine learning can also be applied to discover new control paradigms, see Stanfel et al., 2021 for an example.

1.4 Aim of this work

The articles comprising this thesis are concerned with the development and application of the Lattice Boltzmann Method to the field of wind energy research. Previous work has shown that LBM has a great potential to reduce computational demand for wind turbine simulations in idealized conditions. The work presented here shows how LBM can enable further industrial LES application and provides a validation of the method in realistic conditions. Furthermore, we show that LBM expands the use cases of wind farm LES to rapid generation of training data for machine learning models, to create models to estimate flow fields or train novel control algorithms. Finally, we describe and analyze the physical mechanisms behind the helix approach and how it can be extended to multiple turbines.

1.5 Outline

Part I of this thesis comprises an introduction to the relevant methodologies used in the rest of the thesis and a comprehensive summary of the included articles. Following this introductory chapter, chapter 2 gives an

introduction to the numerical method and selected topics from the field of machine learning. The lattice Boltzmann method is summarized in section 2.1, section 2.2 discusses neural networks and the basics of reinforcement learning are outlined in section 2.3. The summarized findings of the included articles are presented in chapter 3. Concluding remarks and an outlook on future work can be found in chapter 4. Part II consists of the publications included in this thesis.

2. Methodology

The following chapter describes the underlying methods used in this thesis. It begins by describing some fundamental aspects of the lattice Boltzmann method. Furthermore, since papers III and IV are concerned with machine learning and reinforcement learning, a short overview of the respective methodologies is provided as well.

2.1 The Lattice Boltzmann Method

The lattice Boltzmann method originates from the field of kinetic theory. An exhaustive derivation and discussion is provided by Krüger et al., 2017, whose outline I follow in this section. The fluid is described by the distribution of particles in 6 dimensional phase space of space and velocity. The particle distribution function (pdf) $f(\boldsymbol{x}, \boldsymbol{\xi}, t)$ gives the probability of encountering a particle with velocity $\boldsymbol{\xi}$ at time t and location $\boldsymbol{x}$. The distribution in phase space can be related to macroscopic quantities such as density ρ and velocity $\boldsymbol{u}$ through moments of f, that is, integrals of f premultiplied with powers of $\boldsymbol{\xi}$ over velocity space,

$$\rho(\boldsymbol{x}, t) = \int f(\boldsymbol{x}, \boldsymbol{\xi}, t) \mathrm{d}\xi^3 \tag{2.1}$$

$$\rho(\boldsymbol{x}, t)\boldsymbol{u}(\boldsymbol{x}, t) = \int \boldsymbol{\xi} f(\boldsymbol{x}, \boldsymbol{\xi}, t) \mathrm{d}\xi^3. \tag{2.2}$$

Higher order moments give the internal and total energy and further macroscopic quantities.

The Boltzmann equation describes the evolution of the pdf with time,

$$\frac{\mathrm{D}f}{\mathrm{D}t} = \frac{\partial f}{\partial t} + \underbrace{\xi_\alpha \frac{\partial f}{\partial x_\alpha}}_{\text{advection}} + \underbrace{\frac{F_\alpha}{\rho} \frac{\partial f}{\partial \xi_\alpha}}_{\text{external force}} = \Omega, \tag{2.3}$$

where $\frac{\mathrm{D}f}{\mathrm{D}t}$ denotes the total derivative of f and $\boldsymbol{F}$ presents an external body force. The evolution of the pdf is driven by the collision of particles, described by the collision operator Ω. Compared to the Navier-Stokes equation, the structure of the Boltzmann equation is rather short, however, the expression of the collision operator is complex and omitted

here, since it is not relevant for LBM and alternative formulations will be discussed at the end of this section.

Without the presence of external forces, the collision of particles drive the particle distributions towards an equilibrium state, described by the Maxwell distribution function,

$$f^{\mathrm{eq}}(\boldsymbol{x}, \boldsymbol{\xi}, t) = \rho \left(\frac{1}{2\pi RT} \right)^{3/2} \exp\left(-\frac{|\boldsymbol{\xi} - \boldsymbol{u}|^2}{2RT} \right), \tag{2.4}$$

where R is the specific gas constant and T the temperature of the gas. The equilibrium distribution is isotropic around the macroscopic velocity and represents a fluid in a state without viscous stresses.

To use the Boltzmann equation for computational fluid dynamics, it needs to be discretized in time, space, and velocity space. The velocity space is discretized with a small number of discrete velocities, the velocity set. The distribution function is approximated with an Hermite series expansion. Because the generating function of Hermite polynomials has the same form as the Maxwell distribution, only the first three terms of the Hermite series expansion of the distribution function are required to recover the correct macroscopic behavior. Furthermore, by choosing the discrete velocities from roots of the third Hermite polynomial, the Gauss-Hermite quadrature allows the exact computation of moments, which are related to macroscopic quantities, as mentioned earlier. Thus, the values of the equilibrium distribution at the discrete velocities are, in the isothermal case,

$$f^{\mathrm{eq}}(\xi_i) = f_i^{\mathrm{eq}} = w_i \rho \left[1 + \frac{c_{i\alpha} u_\alpha}{c_{\mathrm{s}}^2} + \frac{u_\alpha u_\beta \left(c_{i\alpha} c_{i\beta} - c_{\mathrm{s}}^2 \delta_{\alpha\beta} \right)}{2 c_{\mathrm{s}}^4} \right], \tag{2.5}$$

where w_i is the weight, $c_{i\alpha} = \frac{1}{\sqrt{3}} \xi_{i\alpha}$ are the velocities of the velocity set divided by a $\sqrt{3}$ for convenience and c_{s} is the speed of sound. Note that the rules of Einstein summation convention only apply to greek indices, latin indices denote lattice directions and are exempt from the convention.

Simulations with high Reynolds numbers or rotational symmetries require a highly isotropic lattice, thus a velocity set with 27 discrete velocities is recommended, shown in Figure 2.1. See for example Appendix A in Asmuth, Olivares-Espinosa, and Ivanell, 2020 for a discussion of the effect of a D3Q19 lattice on the simulation of wind turbine wakes.

Time and space are discretized such that particles travel from one grid point to another during one timestep Δt. By restricting the spatial discretization to a uniform cubic grid with constant grid spacing Δx and requiring all lattice speeds to be integer multiples of $\Delta x / \Delta t$, it can be guaranteed that no particle gets stuck between grid point. By applying the method of characteristics, we arrive at the lattice Boltzmann

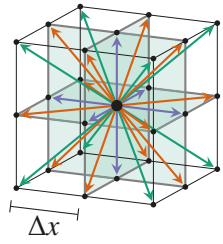

Figure 2.1. Velocity vectors of the D3Q27 velocity set. Colors indicate the length of the velocity vectors and black dots represent nodes of the discretized grid.

equation,

$$f_i(\boldsymbol{x} + \boldsymbol{c}_i \Delta t, t + \Delta t) = f_i(\boldsymbol{x}, t) + \Delta t \Omega_i(\boldsymbol{x}, t). \tag{2.6}$$

The equation is explicit and can be shown to be second order accurate. The computation can be split into two parts, the collision and the streaming step. In the collision step, the post collision pdfs are computed, i.e. the right hand side of (2.6). In the streaming step, the post-collision pdfs are advected to the neighboring nodes. The only component missing at this point is the collision operator Ω_i. Since the aim of LBM is not to recover the correct mesoscopic behavior, there is no need to try to approximate Boltzmann's collision operator. Instead, a computable operator yielding the correct macroscopic behavior has to be found. The first collision operator was proposed by Bhatnagar et al., 1954, usually referred to as BGK operator. It consists of a relaxation of the discrete pdfs towards their respective equilibrium values with a constant relaxation rate τ,

$$\Omega_i = \frac{1}{\tau} \left(f_i^{eq} - f_i \right). \tag{2.7}$$

The relaxation rate is related to the viscosity of the fluid by

$$\nu = c_s^2 \left(\frac{1}{\omega} - \frac{1}{2} \right). \tag{2.8}$$

The BGK operator yields good results at low Reynolds numbers, at higher Reynolds number it is not stable. A lot of research has been conducted regarding the development of new collision operators and analyzing their properties. A thorough analysis of different operators can be found in Coreixas, Chopard, et al., 2019 and Coreixas, Wissocq, et al., 2020. In this work, the cumulant collision operator is used, which was introduced in Geier, Schönherr, et al., 2015 and further developed in Geier, Pasquali, et al., 2017a,b. It is based on a relaxation of the cumulants of distribution function, instead of the pdfs themselves. It exhibits excellent stability, even at high Reynolds numbers and is fourth

order accurate in diffusion. The effectiveness of the method is due to the properties of the cumulants, since they are statistically independent by construction and thus permit relaxation at different rates, as the authors stress in Geier, Pasquali, et al., 2017a.

The widely praised computational efficiency of LBM arises from the just discussed aspects: the simplicity of the scheme due to its explicitness allows an easy implementation, while still being of second order accuracy. Furthermore, the LBM is local in the sense that the computation of the collision operator only requires the pdfs of the local node. Being explicit and local makes the LBM "pleasingly parallel" and thus well suited for modern massively parallel hardware such as GPUs. All non-local operations, i.e. streaming, are linear and reduce to a memory operation on a computer. Furthermore, due to restriction to uniform cubic grids, grid generation is significantly easier than in classical Navier-Stokes based finite volume solvers, even though this might result in larger grids. A major drawback of the method is the difficulty of formulating boundary conditions, since the fundamental quantity is the pdf and not pressure and velocity, thus boundary conditions also have to be formulated in terms of pdfs.

2.1.1 Boundary Conditions

Considerable research effort has been spent on the development of boundary conditions for the LBM. Periodic boundaries are easily implemented, as only the streaming step has to be modified. Solid walls can be modeled with the so-called bounce-back schemes, that also allow setting a constant velocity of the wall. Originally developed for straight walls at the centerline between two nodes, the method has been extended to moving walls and curved boundaries, see for example Bouzidi et al., 2001. Slip boundaries are modeled with a similar approach, the bounce-forward method. A thorough discussion on bounce-back and bounce-forward methods can be found in Krüger et al., 2017, p. 175-189 and 206-208, respectively

A variety of velocity boundary conditions exist, a comparison can be found in Latt et al., 2008, alternatively, the bounce-back approach can be extended, as was already proposed in Ladd, 1994 or Lallemand and Luo, 2003. An alternative approach is proposed in appendix E of Geier, Schönherr, et al., 2015 that only requires local post-collision pdfs.

In order to create correct inflow conditions, wind farm simulations often utilize precursor-successor setups, i.e. a precursor simulation is run as a bi-periodic simulation of the ABL without the wind farm and the velocity sampled in a plane is used as the inflow condition of the successor simulation, with the wind farm present. In LBM, the precursor-successor

scheme can be implemented in several ways. The velocity can be sampled from the precursor and used as velocity for the aforementioned velocity boundary condition, or the necessary pdfs can be sampled and directly used as the missing pdfs at the boundary.

Outflow boundary conditions can be implemented with so-called anti-bounce back approaches. An alternative approach is presented in Appendix F of Geier, Schönherr, et al., 2015. Often the outflow boundary is combined with a damping layer, i.e. an area with increased viscosity to reduce reflected acoustic waves, which is discussed in Xu and Sagaut, 2013.

2.1.2 Large Eddy Simulation

Since it is not feasible to resolve all scales of turbulence in ABL simulations, large eddy simulation is commonly employed. Classical LES is based on applying a filter kernel to the Navier-Stokes equations and then solving the filtered equations while modeling the effect of subgrid-scales in the resolved flow. The same procedure can be applied to the Boltzmann equation, shown for example in Sagaut, 2010. However, in practice, turbulence is modeled according to the Boussinesq hypothesis by computing an effective viscosity $\nu_{\text{eff}} = \nu + \nu_{\text{turb}}$ and the relation (2.8) is used to apply ν_{eff}. Thus the turbulent viscosity ν_{turb} can be computed via all turbulence models developed for Navier-Stokes models.

Turbulence models solely based on the stress tensor, such as the classical Smagorinsky model, proposed in Smagorinsky, 1963, or the QR model proposed in Verstappen, 2011, can be computed from local quantities and are therefore computationally very efficient, while others, such as the anisotropic minimum dissipation model, introduced in Rozema et al., 2015, requiring the full velocity gradient tensor, require the computation of finite differences.

Towards the wall, the turbulent scales become even smaller and can not be appropriately modeled by subgrid-scale models anymore. Therefore, wall models have to be used. The first approach to incorporating wall models in LBM was proposed by Malaspinas and Sagaut, 2014, however, this approach did not turn out to be sufficient at very high Reynolds numbers. Asmuth, Janßen, et al., 2021 proposed an approach to incorporate the wall shear stress predicted by a wall model, called the inverse momentum approach (iMEM), which is stable and accurate at very high Reynolds numbers, bridging the gap to ABL simulations in LBM.

2.1.3 The Actuator Line Method

As was described in section 1.1, the use of actuator methods reduces the computational cost of wind turbine simulations drastically. The original formulation by Sørensen and Shen, 1999 discretizes the blades as lines of points. The velocity is sampled at each point $\boldsymbol{x}_i$ of the line and the relative velocity V_{rel} and the angle of attack α at the blade are computed by

$$V_{\text{rel}} = \sqrt{u_x^2 + \left(u_y + \omega r\right)^2} \tag{2.9}$$

$$\alpha = \arctan\left(\frac{u_x}{u_y + \omega r}\right) - \theta - \phi, \tag{2.10}$$

where u_x and u_y are the downstream and tangential velocity in the static blade frame, r is the radial position of the blade section, θ is the pitch of the blade and ϕ the local twist angle. The angles, velocities, and forces are shown in Figure 2.2. The lift and drag coefficients of the blade airfoil are computed from a lookup table and the line force exerted by the blade section $\boldsymbol{f} = [f_n, f_t]^T$ is computed as

$$\boldsymbol{f} = -\frac{1}{2}\rho c V_{rel}^2 \left(C_L(\alpha)\boldsymbol{e}_L + C_D(\alpha)\boldsymbol{e}_D\right) \tag{2.11}$$

where c is the chord length of the airfoil and $\boldsymbol{e}_L$ and $\boldsymbol{e}_D$ are unit vectors in the direction of lift and drag, respectively. The forces are then integrated along the blade sections and applied as body forces in the fluid domain. As mentioned previously, body forces can be applied directly in LBM. To avoid numerical instability, the body forces are smeared in the domain, typically with a Gaussian kernel,

$$\eta(r)_\epsilon = \frac{1}{\left(\epsilon\sqrt{\pi}\right)^3} \exp\left(-\left(\frac{r}{\epsilon}\right)^2\right), \tag{2.12}$$

with ϵ being a parameter, the so-called smearing width. Troldborg, 2009 recommends a value of $\epsilon = 2\Delta x$ to avoid numerical instabilities in

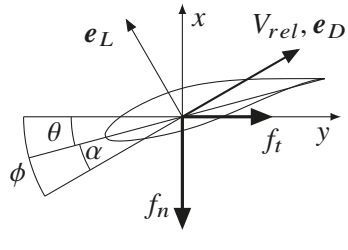

Figure 2.2. Angles, velocities and forces in the local blade frame at the airfoil sections. Forces are shown acting on the fluid.

Navier-Stokes based solvers, which was confirmed by Asmuth, Olivares-Espinosa, Nilsson, et al., 2019 to be suitable for LBM as well. The body force $\boldsymbol{F}$ acting on an LBM node is then the result of a convolution of the line forces with the smearing kernel,

$$\boldsymbol{F} = \sum_{i=1}^{N_b} \int_0^R \boldsymbol{f}(r)\eta(|\boldsymbol{x} - r\boldsymbol{e}_i|)\mathrm{d}r, \tag{2.13}$$

where N_b is the number of blades, R is the radius of the blade and $\boldsymbol{e}_i$ is the unit vector pointing along the blade.

The smearing of the body forces with the Gaussian kernel introduces an error. Meyer Forsting et al., 2019 show that the vortex resulting from the actuator line has the form of a Lamb-Oseen vortex, with a viscous core, while lifting line theory assumes an inviscid core. This results in an overprediction of forces near the blade tip, for example observed in Asmuth, Olivares-Espinosa, and Ivanell, 2020. A growing number of corrections for this error are being proposed, for example in Martínez-Tossas and Meneveau, 2019, Meyer Forsting et al., 2019 and its simplification in Meyer Forsting et al., 2020, Dağ and Sørensen, 2020 and, most recently, Kleine et al., 2023. As of now, no comparison of these models has been conducted in terms of performance and computational efficiency.

2.2 Artificial Neural Networks

Artificial neural networks (ANNs) are the basis of a wide range of machine learning algorithms. ANNs comprise several layers, which can be of different kinds, each taking the output from the previous layer as the input. A feed-forward layer consists of several neurons with individual weights. The k-th layer of the network takes an input z^k, i.e. the output from the previous layer, applies a weight matrix $\boldsymbol{W}^k$ and a bias vector $\boldsymbol{b}^k$ and finally an activation function f to compute the new output y_j^k,

$$y_j^k = f(w_{ij}^k z_i^k + b_j^k). \tag{2.14}$$

Further expansion on neural networks can be found in Demuth et al., 2014. While feed-forward networks are very simple and flexible, they can also be difficult to train and grow rapidly in size. They also do not possess any kind of memory.

Convolutional neural networks (CNNs) were developed to deal with pixelated data, originally images, by Fukushima, 1980. Each pixel has several channels. A convolution layer takes an input matrix and then applies a kernel to a sub-matrix, as is depicted in Figure 2.3, therefore

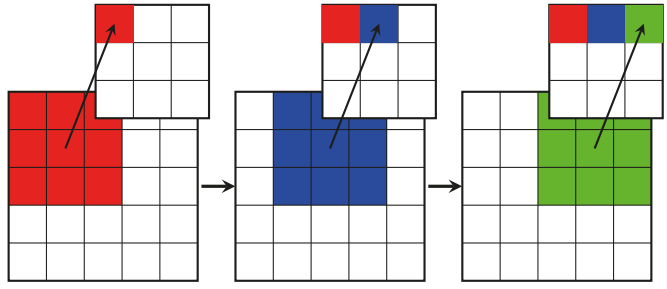

Figure 2.3. Schematic of a convolutional layer.

retaining geometric information. A more formal definition can be found in Goodfellow et al., 2016, ch. 9. The kernel is similar to a layer in a feed-forward network, since it also applies a weight matrix and bias to the input matrix and then an activation function. The number of output channels can differ from the number of input channels. CNNs are applicable to much larger input data, since they only have to learn the parameters of the kernel, whereas a feed-forward network has to learn weights for each individual point. The inverse operation to the convolutional layer is the transposed convolutional layer. A classical architecture for CNNs is the encoder-decoder architecture. In the encoder, the size of the data is successively reduced by applying convolution layers in combination with MaxPooling layers that reduce the size, while increasing the number of channels, thus encoding more features. The opposite is done in the decoder, i.e. transposed convolutions are applied, decreasing the number of channels while increasing the size of the output.

Another type of network is the recurrent neural network (RNN). RNNs can retain information between inputs and are often used for temporal data. A common type of neuron used in RNNs is the long short-term memory (LSTM) cell, which was proposed by Hochreiter and Schmidhuber, 1997. Its inputs consist of its output from the previous timestep, an internal cell state, and the output from the previous layer. The cell comprises a forget, an input and an output gate.

Regardless of the type of network, the training process consists of incrementally adjusting the set of weights $\boldsymbol{\Theta}$ of the network. Weights are updated based on some optimization algorithm, such as stochastic gradient descent (SGD) or one of its extensions such as the Adams algorithm, introduced by Kingma and Ba, 2017, in order to maximize an objective function J. SGD computes updates the weights from iteration n to $n+1$ according to

$$\boldsymbol{\Theta}_{n+1} = \boldsymbol{\Theta}_n + \alpha \nabla_{\Theta} \mathrm{J}_{\Theta}, \tag{2.15}$$

where the learning rate α controls how large the updates are. The optimization depends on computing $\nabla_\Theta \mathrm{J}$, which can be efficiently computed via a method known as backpropagation.

2.3 Reinforcement Learning

Reinforcement learning imitates the trial-and-error mechanism which humans use to learn how to interact with their environment. The basics I present here can be found in Sutton and Barto, 2018. The underlying model is the Markov Decision Process which models interactions between an agent and the environment. The agent performs an action, $\boldsymbol{A}$, and the environment responds with a new state, $\boldsymbol{S}$, and a reward, R. The action is chosen based on a policy π. The goal of the learning process is now to find a policy such that the long-term reward is maximized. To do so, the interactions are broken up into episodes with a finite number of interactions. The long-term reward of timestep t is known as the return G_t, typically computed as

$$\mathrm{G}_t = \sum_{t'=t}^{T} \gamma^{t'-t} \mathrm{R}_{t'}, \tag{2.16}$$

where T is the number of timesteps in the episode and γ is the discount rate. A low discount rate emphasizes the long-term return while higher discount rates favor short-term rewards.

Different ways of determining the policy exist. In policy gradient methods, the policy function π is optimized directly. The policy is formulated as a probability distribution $\pi_\Theta(\boldsymbol{A}|\boldsymbol{S})$, with a set of parameters $\boldsymbol{\Theta}$ that are optimized in the learning process.

During one training episode, the action $\boldsymbol{A}$ is sampled from $\pi_\Theta(\boldsymbol{A}|\boldsymbol{S})$ and $\boldsymbol{A}_t$, $\boldsymbol{S}_t$ and R_t are recorded. After the episode is completed, G_t can be computed and $\boldsymbol{\Theta}$ is updated to maximize some performance measure J. As described in section 2.2, the update can be performed with SGD or similar optimization algorithms, which requires the computation of $\nabla_\Theta \mathrm{J}(\boldsymbol{\Theta})$. In its simplest form, $\nabla_\Theta \mathrm{J}(\boldsymbol{\Theta})$ can be approximated as

$$\nabla_\Theta \mathrm{J}(\boldsymbol{\Theta}) \approx \frac{1}{N_e} \sum^{N_e} \frac{1}{T} \sum_{t}^{T} \mathrm{G}_t \frac{\nabla_\Theta \pi_\Theta(\boldsymbol{A}_t|\boldsymbol{S}_t)}{\pi_\Theta(\boldsymbol{A}_t|\boldsymbol{S}_t)}, \tag{2.17}$$

where N_e is the number of episodes over which is averaged. This algorithm is named the REINFORCE algorithm, introduced in Williams, 1992, however, it is very inefficient and requires a lot of training time.

The proximal policy optimization (PPO), proposed by Schulman et al., 2017, is a performance measure that approximates the improvement

in return due to a policy update, greatly reducing the training time. Two functions have to be found in PPO, the policy and a function estimating the performance improvement due to the policy update. Both functions are typically parameterized as neural networks, although this is not required. Parameterization through a neural network allows the use of backpropagation to compute the necessary derivatives in (2.17) and the more advanced performance measures.

3. Summary

This thesis is composed of papers considering the development and application of LBM to wind farm modeling and wind farm control and its usage in an industrial context. The first paper establishes what conditions have to be met by an LES tool to be considered for industrial application and what use cases could look like, while also comparing runtime requirements to the current state of the art of LBM simulations. Paper II consists of a validation study against measurements and provides a comparison in accuracy and performance to established LES tools. Papers III and IV use LBM to explore different machine learning techniques for wind farm modeling and wind farm control, respectively. Finally, paper V provides a thorough analytical and numerical examination of the helix approach, one of the wind farm control approaches tested in paper IV.

3.1 Industry perspectives on the application of LES

Engineering models remain the most widely used tools in the industrial sector, while current and future modeling challenges can only be solved reliably through the use of LES, as discussed, for example, in Veers et al., 2019. These challenges include the impact of larger and more flexible rotors, complex terrain, blockage, deep-array effects, and many more. However, the exceptional computational cost and runtime limit the use of LES, not only in the wind energy sector but many industrial applications of LES. The gap between research, where LES has been in wide use for well over 20 years, and the industrial application, was named the "LES-crisis" by Löhner, 2019. Löhner et al. postulate that the ability to run a simulation over night is the threshold, below which industrial use of a new method becomes common practice. LBM can reduce runtime of wind farm simulations by multiple orders of magnitude, as was shown by Asmuth, Olivares-Espinosa, Nilsson, et al., 2019, yet the question remains, is it fast enough?

To assess what use cases industrial experts see for LES, what they see as challenges, and what requirements would have to be fulfilled for them to consider using LES, an open survey was conducted. The survey was distributed online by posting it on social media sites and in email newsletters as well as through direct contacts. A total of 48 industry

experts responded. It is not possible to say whether this is a representative overview of the industry, however, random subsampling of half of the respondents yielded similar statistics. To assess the quality of the

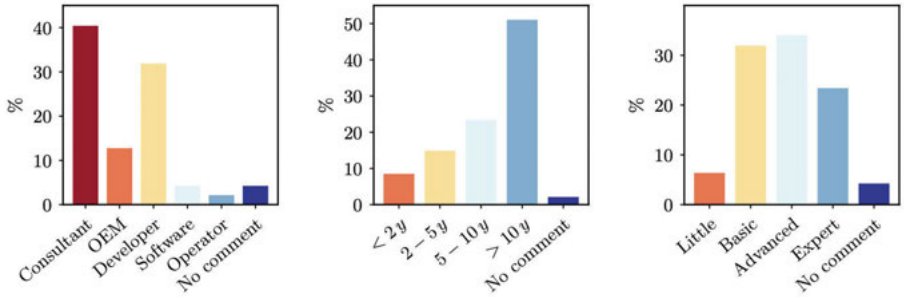

Figure 3.1. Overview over the background of the survey responds. (left) Main field of activity of the company, (middle) number of years of experience with CFD and (right) level of expertise, where "basic" refers to Batchelor or Masters level course in CFD, "advanced" means several years of actively using CFD and "expert" refers to having actively developed CFD or a PhD in CFD. Figure adapted from paper I

responses, the respondents were asked which type of company they work for and the amount of experience they have, measured in number of years and level of expertise. The results, shown in Figure 3.1, confirm that all stakeholders of the industry are represented and that the respondents have a high level of expertise.

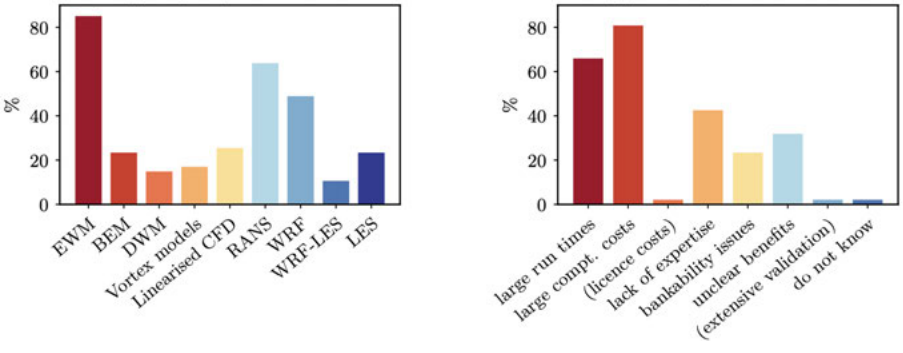

Figure 3.2. Model practices in the companies of the respondents. (left) regularly used types of models, (right) main factors limiting the use of CFD. Answers in parentheses from free answers. Figure adapted from paper I.

The participants were also asked to indicate which type of models they most commonly use, the results are shown on the left in Figure 3.2. Engineering wake models were the most commonly used methods, while mid-fidelity methods, such as dynamic wake meandering models and linearized CFD methods, were only used by a fifth or less of the partici-

pants. Reynolds averaged Navier Stokes (RANS) methods were used by more than half of all participants. About 20 % of respondents indicated to regularly use LES. When asked about the factors limiting the use of LES, larger runtimes and computational cost were selected as the most common reasons, as shown in Figure 3.2 on the right, which is in line with the common notion that LES is too expensive. However, 40 % of respondents also cited lack of expertise as a reason, about 30 % did not see clear benefits in using LES and about 20 % referred to bankability issues. While those reasons can not be addressed directly through further academic development of the method, academia does play a role in training personnel and developing expertise, but also proving where LES can provide more accurate results compared to other methods in order to develop trust in the methodology. An increased trust in the accuracy of results will also result in better bankability and proven improvements in accuracy might also elucidate in which cases LES is beneficial.

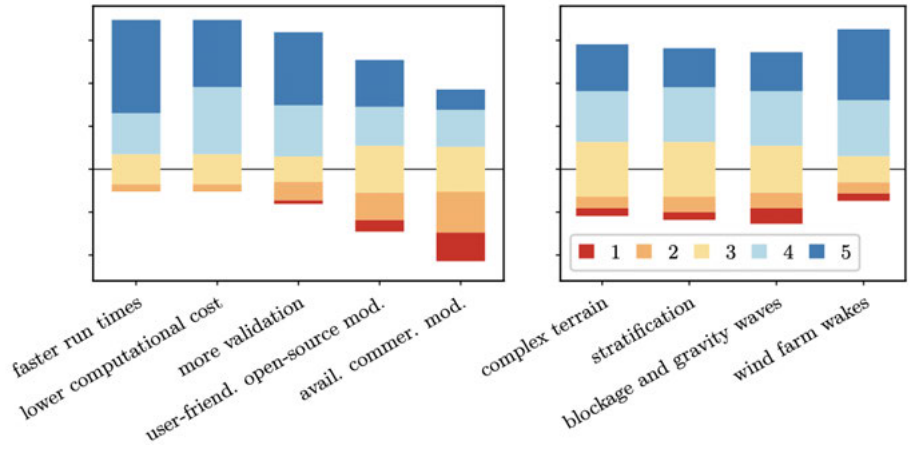

Figure 3.3. (left) Importance of different requirements an LES model has to fulfill in order for the model to be considered. Ranked from 1, "not important at all/the status quo is sufficient" to 5, "absolutely necessary/requires major improvements". (right) modeling challenges that would benefit from the use LES, ranked on a scale 1 to 5, with 1 referring to "not relevant at all/low-fidelity models are sufficient" and 5 being "high relevance/lower-fidelity models are inadequate". Results shown as diverging stacked bar charts (total length referring to 100 %) centered around the middle of category 3. Ticks on the vertical axis mark intervals of 25%. Figure adapted from paper I.

The left panel in Figure 3.3 shows what specific requirements an LES model has to fulfill for the respondents to use such a model. Again, faster runtimes and lower computational cost were indicated as the primary requirements. However, more validation was also ranked important or very important by the vast majority of respondents. Interestingly, user

friendly open source models were seen as more important than commercial models. The three requirements ranked most important indicate a clear direction for further research, if academia wants to facilitate the application of LES in the industry.

In order to determine areas on which to focus model development, the respondents were asked for which phenomena they see the largest benefits of using LES. The results are shown in Figure 3.3 on the right. Wind farm wakes were marked important or very important by a large majority, but respondents reported interest in all other phenomena as well, indicating an interest in LES by many industrial users.

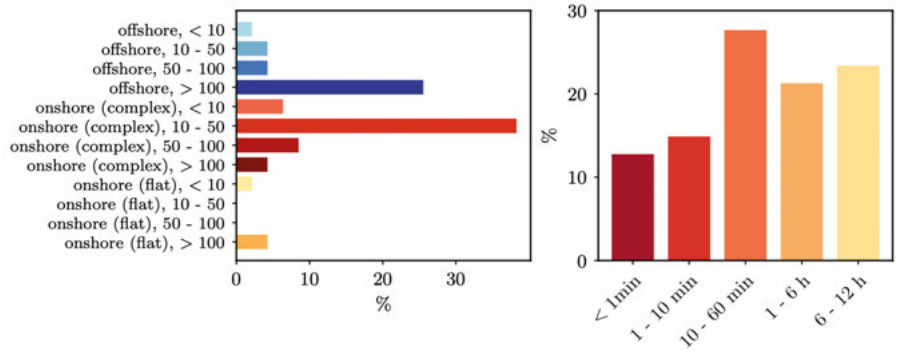

Figure 3.4. Specification of a potential use case for LES. (left) Number of turbines and classification of terrain. (right) Maximum runtimes acceptable for a simulation of 1 hour. Figure adapted from paper I.

Finally, the respondents were asked to define their typical use case and what time requirements would have to be met by an LES tool for the respondents to consider LES. The responses are shown in Figure 3.4. Two typical use cases emerge as the most common application, namely very large offshore wind farms with more than a hundred turbines and onshore wind farms with 10 to 50 turbines. These cases also reflect the areas in which the respondents expected the largest benefits from the use of LES. The reported runtime requirements vary, with more than 10 % of respondents requiring less than one minute runtimes for 60 minutes of simulation while also over 20 % indicating acceptable runtimes on the order of 6 to 12 hours.

Finally, we conducted an exemplary simulation of 64 DTU 10 MW reference turbines (Bak et al., 2012) with a diameter of 198 m. The domain had a size of 20 km × 15 km × 1.6 km in streamwise, spanwise and vertical direction and an isotropic grid spacing of 10 m. Towards the top, the domain was coarsened two times. The simulation had a timestep of

$\Delta t \approx 0.1\,\mathrm{s}$[1] in the most refined region and statistics were computed for 6000 s after an initial spin-up time of 2000 s. The inflow boundary was obtained from a precursor simulation, while Monin Obukhov similarity theory was used to compute a wall stress applied with iMEM at the ground. Further details of the simulation can be found in the original publication. The simulation was carried out on a single NVIDIA A100 GPU. The total wall time of the simulation amounted to 5112 s, thus the simulation was performed at a ratio of simulated to wall time of 1.54. This computational performance meets the requirements reported by about 45 % of the respondents.

In summary, we showed that experts from across the industry see LES as a vital tool to address current and future modeling challenges. Along with further development of the method, trained personnel and additional validation are important for improved adoption by the industry. However, the current state of the art of LBM simulations is already capable of achieving the runtimes required by a large share of the respondents.

3.2 Validation against full scale measurements

Paper I mentions that one of the reasons limiting the adoption of LES for industrial use is a lack of validation. Validation is a crucial step to build trust in models and to find strengths and weaknesses of different methods. Several cross-code comparisons and validation studies of wind turbine simulations have been conducted over recent years, such as Martínez-Tossas, Churchfield, et al., 2018, Lignarolo et al., 2016, and Asmuth, Navarro Diaz, et al., 2022. In paper II, we conduct a validation against measurements of a full scale turbine. The measurements were published in the course of the SWiFT (Scaled Wind Farm Technology) benchmark by Doubrawa, Quon, et al., 2020, comparing 13 models of different fidelity, including five LES models. The examined turbine is a Vestas V27 turbine with a rotor diameter of $D = 27\,\mathrm{m}$ and a hub height of 32.1 m. The benchmark data consists of measurements by a meteorological tower $2.5D$ upstream of the turbine, torque, rotor speed and power of the turbine, and a SpinnarLidar (T. Mikkelsen et al., 2013) pointed towards the wake, more details can be found in Doubrawa, Debnath, et al., 2019. Benchmarks in stable, neutral, and unstable conditions were published, however, in the present study only the neutral case is considered. Each benchmark dataset contains five windows of 10 minutes with minimal synoptic forcing.

[1] Corrigendum to paper I: The original publication contains a misprint. A placeholder for the timestep size was not removed before publication. Therefore I mention the correct timestep size here.

The benchmark data, results from the participating LES simulations as well as results from another validation study by Jézéquel et al., 2021 are compared to results from a single simulation conducted with the LBM solver `VirtualFluids` coupled to the wind farm coupling module `WiFI` mainly developed as part of this project. The simulation utilizes a precursor-successor setup. The precursor has a domain of 6 km × 6 km × 1 km in streamwise, lateral and vertical direction with an isotropic grid spacing of 7.8125 m. The boundaries are periodic in streamwise and lateral direction and a free-slip condition is applied at the top while the wall stress prescribed from Monin-Obukhov similarity theory is applied with an iMEM boundary at the bottom. The flow is driven with a constant pressure gradient tuned to match the friction velocity of the benchmark data. Further details of the simulation can be found in the original publication in II. The successor simulation had the same outer dimensions as the precursor but refinement zones around the turbine which is placed at (3000 m, 2500 m, 32 m). The grid spacing in the finest region around the turbine is about 1 m. Further details can be found in the original publication. The turbine is modeled with the actuator line method, employing the smearing correction proposed by Meyer Forsting et al., 2020.

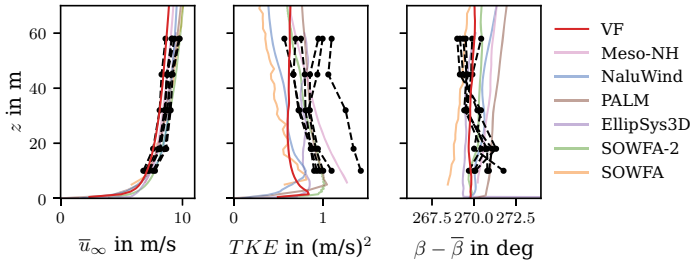

Figure 3.5. Inflow profiles. (left) Mean inflow velocity, (center) turbulent kinetic energy, (right) inflow direction. Black dashed lines are the averages from benchmark measurements. Figure adapted from paper II.

Figure 3.5 shows the comparison of the results from the simulation to the other participants and the measurement data. The vertical mean velocity profile, shown in Figure 3.5 on the left, matches the profile of the measurements and the other simulations, albeit that the mean velocity is lower than in most other results. This lower velocity was due to streaks of lower and higher velocity observed in the measurements. We chose to place the turbine in a streak of lower velocity since the velocity was constant throughout the simulation at the cost of not matching the mean velocity as well. The turbulent kinetic energy, shown in Figure 3.5 in the middle, is also lower than in most other simulations and in the measurements. We aimed to match the turbulence intensity at hub height, but

due to the lower inflow velocity, we also had to aim for a lower turbulence intensity. Doubrawa, Quon, et al., 2020 report a turbulence intensity of 10.7 % at hub height while our simulation has a turbulence intensity of 9.7 %. The results shown in Figure 3.5 on the right indicate that no veer was present in the simulation and therefore no streaks passed through the observed part of the domain.

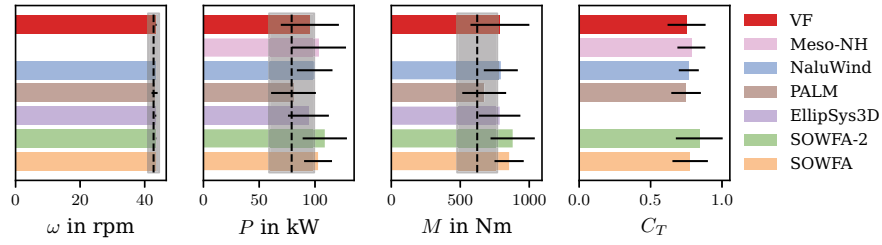

Figure 3.6. Turbine response. Average rotor speed, average power production, average generator torque and average thrust coefficient from left to right. Black solid lines show one standard deviation above and below the average. Grey shaded area represents two standard deviations around the average obtained from the ensemble average, shown in black dashed lines. Figure adapted from paper II.

The turbine response is shown in Figure 3.6. All quantities for which measurements are provided are matched well, while most other models overpredict torque and power, a common problem in actuator line simulations due to the smearing error discussed in subsection 2.1.3. The thrust coefficient was not measured, but our results compare well to the other simulations, despite a rather low resolution in our simulation. We attribute this mainly to the effect of the smearing correction, which allows accurate computation of the blade forces at lower resolution, as shown by Meyer Forsting et al., 2020.

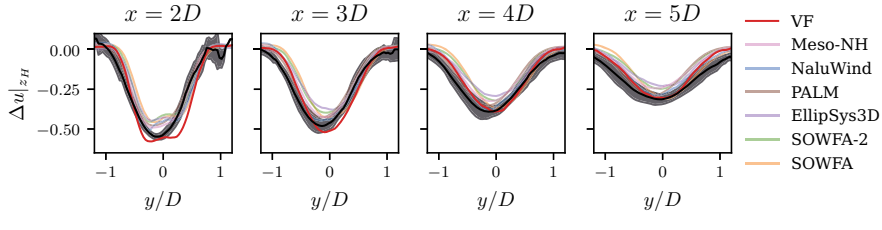

Figure 3.7. Velocity deficit profiles at hub height. Black line represents the ensemble average and grey shaded area the ensemble standard deviation. Figure adapted from paper II.

The mean velocity deficit profiles at hub height at four distances downstream of the turbine are shown in Figure 3.7. At a distance of $2D$

downstream of the turbine, the measured profiles hava a single peak while our simulation results still has two peaks, indicating that the transition from near to far wake has not yet been completed in the simulation. It is also noteworthy that the other simulations report velocity deficits weaker than the measurements, while our simulations exhibits stronger velocity deficits. Further downstream, our results are very close to the measurements, while most other simulations report weaker velocity deficits. Overall, all LES simulations captured the correct shape and trends, while the strength of the wake deficit are underestimated by the other participants and overestimated in our results.

The wake strength and recovery are also examined by comparing the maximum velocity deficit and the difference of the maximum velocity deficit to the maximum velocity deficit at the first measurement distance, i.e. $2D$. Both quantities matched well to the measurements, I omitted the results here for brevity's sake.

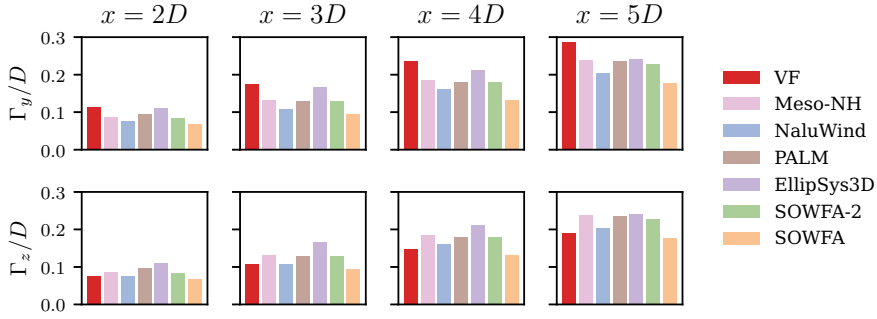

Figure 3.8. Horizontal (top) and vertical (bottom) wake meandering. Figure adapted from paper II.

Furthermore, the wake meandering is compared with the other simulation results in Figure 3.8. Measurement data is not available. The lateral wake meandering is stronger in our simulation than in all other simulations, while the vertical wake meandering is weaker. Overall, the results vary between all the simulations and as no measurements are available it is not possible to differentiate which results are more realistic. The setup between the simulations is quite different, for example the domain size, yet no trend emerged, e.g. that larger simulations with larger structures report larger wake meandering.

Finally, a comparison of the computational performance of the different methods to the accuracy of the simulation for all participants is shown in Figure 3.9. The computational performance is measured in energy to solution ($E2S$), which is the total energy consumed by the hardware to arrive at a solution. This quantity makes it possible to compare heterogenous hardware, and is also less sensitive to the number

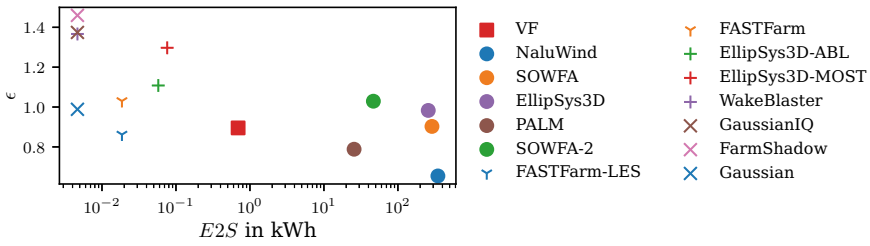

Figure 3.9. Accuracy compared to energy to solution for all models that participated in the benchmark. Circles represent LES models, three-pointed star are dynamic wake meandering models, crosses are RANS models, and xs are engineering wake models. Figure adapted from paper II.

of processors used in a homogeneous system. Only the number of core hours was reported in the original publication, so we converted the core hours to $E2S$ by assuming a modern CPU. The accuracy ϵ is measured as the average root mean squared error of the horizontal and vertical velocity deficit of the simulation of the ensemble average of the measurements. The energy consumption reported for the traditional LES already displays a large spread, varying by more than one order of magnitude. However, our simulation is more than one order of magnitude more efficient than the most efficient traditional LES simulation and even more than two orders of magnitude more efficient than the majority of the LES simulations. Meanwhile, our simulation ranged in the middle in terms of accuracy. Nevertheless, it is also still an order of magnitude more expensive than the next lower fidelity, i.e. RANS. However, the RANS simulations also reported a higher error.

Overall, the results show that our methodology is on par in terms of accuracy with established solvers, with strengths and weaknesses. However, the computational efficiency sets it apart from the traditional methods. As mentioned in the previous section, validation is very important for acceptance by industrial users but of course also for researchers. While we successfully showed that we can model a small turbine with decent accuracy, validation against a more modern, much larger turbine is still missing. This is mainly due to a lack of publicly available validation data. Furthermore, as was presented in paper I, the phenomena for which LES is hoped to be the most beneficial, i.e. complex terrain and farm effects, were also not present in this study and thus could not yet be validated.

3.3 WakeNet 0.1 - A convolutional neural network for wind turbine simulations

The previous sections have discussed industrial requirements for the use of LBM and shown that accurate simulations are possible while drastically reducing the computational demand. The reduction in computational also enables new applications of high-fidelity methods, for example to generate highly accurate training data for machine learning models. I have discussed the basics of convolutional neural networks in section 2.2 and how they are useful for pixelated data. Paper III is concerned with the training of a CNN to predict the average fluid field in the wake of a single turbine based on data generated by LBM.

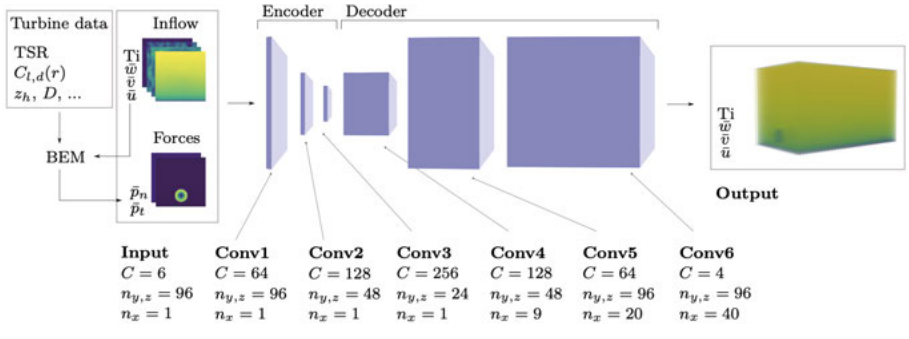

Figure 3.10. Architecture of WakeNet. Figure adapted from paper III.

The model takes the average velocities, $\overline{u}, \overline{v}, \overline{w}$, and turbulence intensity TI at the inflow plane as input, as well as the radial distribution of lift and drag of the rotor, $C_{l,d}(r)$, the hub height z_h, the rotor diameter D, and the tip-speed ratio TSR. The turbine data and tip-speed ratio are fed into a BEM code, as described in M. O. L. Hansen, 2008, that calculates the normal and tangential forces in the plane, p_n and p_t, respectively. The force, velocity and turbulence intensity fields are used as channels for the input to the convolutional neural network. The neural network has an encoder-decoder structure, where the encoder consists of layers of 2D convolutions, while the decoder consists of transposed 3D convolutions, thus extruding the wake velocity field in the downstream direction. The output of the neural network is the average velocity and turbulence intensity in 40 planes with a distance of $D/4$. The model is trained to minimize the mean-squared error of all output quantities. A schematic of the whole model is shown in Figure 3.10, and further details on the layers can be found in the original publication. One key feature of this model is that one output plane can be used as the input plane for the model. That way, many instances of the model could potentially be chained together to model a whole wind farm.

To train the model, a set of 900 simulations of a single turbine are conducted. The simulated turbine is an NREL 5MW turbine (Jonkman et al., 2009) for its definition. Five parameters, namely TSR, hub height, inflow velocity, roughness length, and turbulence intensity, were varied to obtain a total of 900 distinct simulations. The simulations are conducted with the cumulant LBM solver `elbe` and subgrid-scales are modeled with the anisotropic minimum dissipation model, see Rozema et al., 2015. The inflow velocity is constructed from a superposition of a mean log-law velocity and synthetic turbulence generated with the model by Mann, 1998. The lateral boundary is periodic and iMEM is used at the bottom boundary to apply a wall shear stress. The simulation is spun up for one domain flow through time and statistics are then gathered for another ten flow-through times. Each case takes about 0.45 h to complete on a single NVIDIA RTX 2080 TI. Further details can be found in the original publication.

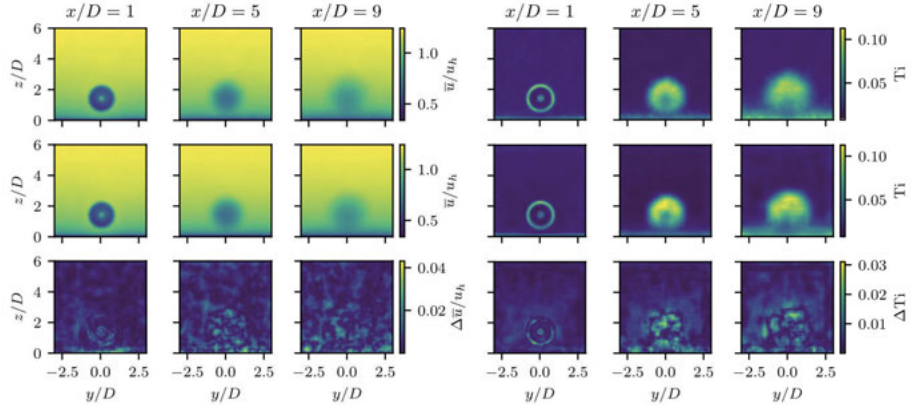

Figure 3.11. Comparison of simulation results (top), prediction from WakeNet (middle) and difference (bottom) of mean velocity (left) and turbulence intensity (right) for one specific case. Figure adapted from paper III.

80 % of the dataset is used as training data while the other 20 % is used as a test set. After training, the model predicts the streamwise velocity component with an $L2$ error of 0.56 %, while lateral and vertical velocity have errors of 27.95 % and 27.41 % respectively. The error in turbulence intensity is 17.01 %. An example comparing the simulation data to the network's prediction is shown in Figure 3.11. The general features of the wake are well captured, including the shape and decay of the wake. The maximum error in the streamwise velocity component stays well below 5% of the hub height velocity u_h, while the error in turbulence intensity reaches 30 %. Nevertheless, the prediction capabilities of the model compare favorably to other comparisons between different methods, as shown for example in Figure 3.7 in the previous section.

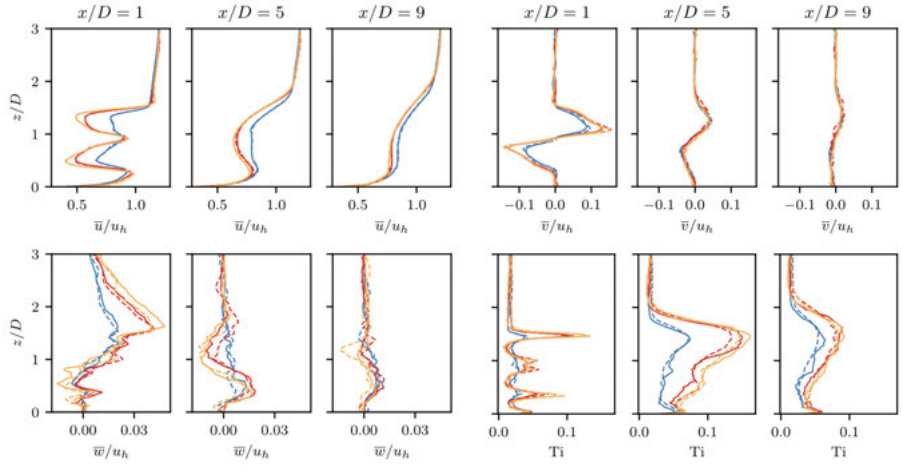

Figure 3.12. Vertical profiles for streamwise (top left), lateral (top right), vertical (bottom left) velocity and turbulence intensity (bottom right) of WakeNet prediction (solid lines) and simulation (dashed lines). Colors refer to different tip speed ratios. Figure adapted from paper III.

A look at the vertical profiles for three different tip-speed ratios, presented in Figure 3.12, shows that the model accurately predicts the streamwise velocity component for each case, thus successfully encoding the effect of changing thrust on the wake. Also effects of transition, such as the change in the shape of the velocity deficit or the increase of turbulence intensity after transition are captured well. The relative error in the lateral and vertical velocity components is larger, the absolute error however is comparable to that of the streamwise velocity.

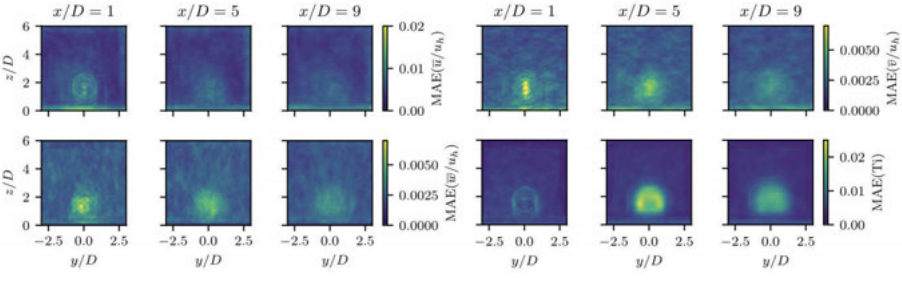

Figure 3.13. Local mean absolute error across the whole test set of average streamwise (top left), lateral (top right) and vertical velocity (bottom left) and turbulence intensity (bottom right). Figure adapted from paper III.

Finally, the mean absolute error (MAE) for each point is depicted in Figure 3.13. The highest error in the streamwise component can be found close to the turbine in the rotor-swept area, but the error vanishes further downstream. The error in lateral and vertical directions behaves

similarly. The error in turbulence intensity is highest in the mid-distance plane, closest to the point of transition. The magnitude in Ti varies the most in the region between $x = 3D$ to $8D$ due to different transition points. Consequently, the model performs worst in this area.

While the generation of training data and training the model requires considerable computational resources, the model executes with an average of $33\,\mu$s on a single NVIDIA RTX 2080 Ti, thus it can provide highly accurate predictions at the computational cost of an algebraic wake model. As mentioned previously, the design of the model allows to chain many instances of the model together to simulate an entire wind farm, which has been done with other machine learning methods for wind farms, for example by Bove et al., 2023. However, the training data does not contain waked inflows, which would increase the parameter space and therefore the number of cases in the dataset considerably. This study shows how the high computational efficiency of LBM enables the generation of training data to train even faster machine learning methods. Of course, these methods lack the physical basis and accuracy of LES, yet they still provide higher accuracy than many engineering models currently in practice at about the same computational cost.

3.4 Discovering New Control Strategies with Reinforcement Learning

Reinforcement learning has had great success in many difficult control problems, most famously beating human players at a variety of strategy games, see Silver et al., 2018, but also for example in the control of drones Gaudet et al., 2020. The application of reinforcement learning to fluid dynamics has garnered a lot of interest recently, after an initial study by Rabault, Kuchta, et al., 2019 proved the potential of the method. However, the studied flow cases are highly simplified and only two-dimensional. A recent overview of reinforcement learning applications in fluid dynamics can be found in Vignon et al., 2023.

Wind farm control has been studied by the wind energy community for a long time and a variety of machine learning methods have been applied, however, most of them with the aim of obtaining optimal static setpoints, see, for example, Park et al., 2017 or Serrano González et al., 2015. On the other hand, dynamic approaches to wind farm control have shown also very promising results, such as Goit and Meyers, 2015 and the helix approach by Frederik et al., 2020b. In the helix approach, the blades of the turbines are pitched sinusoidally, which leads to a helical wake and enhanced wake recovery. I will discuss it in more depth in section 3.5. In the time since the publication of paper IV, other authors have also applied reinforcement learning to optimize different control approaches,

for example Stanfel et al., 2021 and Liew et al., 2023. However, no other publication currently available conducted training in a high-fidelity environment, to the best of my knowledge.

The goal of paper IV is to train a reinforcement learning agent in an environment consisting of three NREL 5MW turbines in a row. Agents control either the pitch amplitude of the helix approach or the generator torque.

Proximal policy optimization is used to train the agent. The policy is parameterized as a recurrent neural network using long short-term memory cells, as introduced by Hochreiter and Schmidhuber, 1997. Further details on the neural network are found in the original publication. The state of the environment consists of point probes in the domain providing instantaneous streamwise velocity as well as the rotor speed of the turbine. The action is either the pitch amplitude or the generator torque and the total generator power is used as the reward. The policy is optimized using the Adam optimizer.

The simulation domain has a size of $6\,\mathrm{D} \times 6\,\mathrm{D} \times 19\,\mathrm{D}$ in streamwise, lateral, and vertical direction. The turbines have a spacing of $5\boldsymbol{D}$. Mann synthetic turbulence with a turbulence intensity of 5% is superimposed on a uniform inflow velocity of 10.5 m/s. The velocity probes constituting the state of the environment are placed in groups of 13 in the shape of a cross with a span of 1.5 diameters in cross-stream planes. Two crosses are placed upstream of the first turbine and four are placed after the first and second turbine each.

In total three agents are trained, one controlling the pitch amplitude of the helix approach, referred to as Agent H, and two agents controlling the generator torque, referred to as M-short and M-long. Agents H and M-short both emphasize short-term rewards, while agent M-long emphasizes long-term rewards. The specifics can be found in the original publication. A reference simulation is conducted, in which all turbines are run with greedy control, i.e. at their optimal tip-speed ratio.

The results of training agent H are shown in Figure 3.14, which shows the action taken by the agent, i.e. the set amplitude for the pitching at the top and the individual and total power on the bottom. The agent is trained for about 3750 flow-through times, corresponding to over $5 \times 10^6\,\mathrm{s}$ of simulated time. The total power exhibits a steady, yet slow increase throughout the training period. The applied pitch angles at turbines 1 and 2 decrease through most of the training time while the pitch angle at turbine 3 does not show as clear of a trend. Due to the helix approach, the power production by turbine 2 is always significantly higher than in the reference simulation. Decreasing the pitch angle of turbine 1 does not appear to reduce the effectiveness of the helix approach while increasing the power production by turbine 1. A decrease in power production by turbine 2 is only observable towards the end of

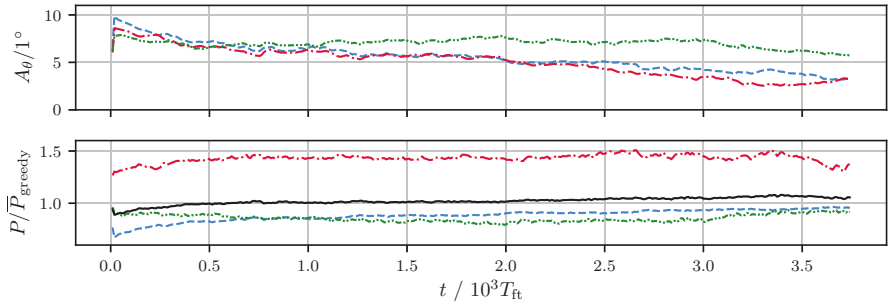

Figure 3.14. Total, Turbine 1, Turbine 2, Turbine 3. Rolling average of pitch amplitude (top) and power production (bottom) of the wind farm throughout the training of agent H. Respective powers normalized by mean per turbine of the reference simulation. Figure adapted from paper IV.

the training period. Note that the turbine power is normalized by the respective average power of the reference case, however, the power production by turbine 1 is significantly higher than by turbines 2 and 3. Thus optimizing the power production by turbine 1 leads to higher improvements in reward. On the other hand, optimizing the power output by turbine 3 yields only little improvement in reward.

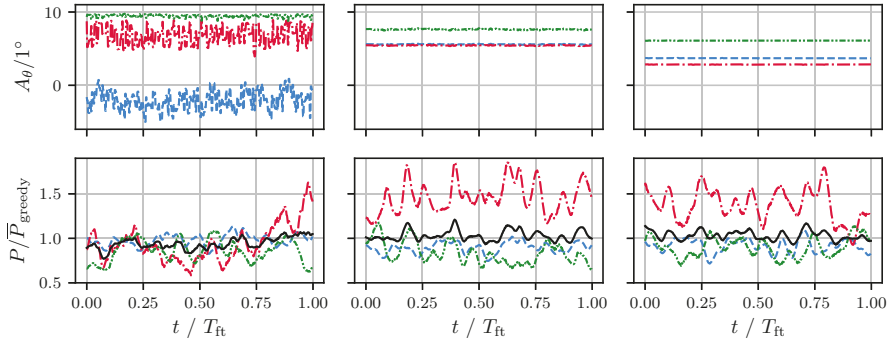

Figure 3.15. Total, Turbine 1, Turbine 2, Turbine 3. Amplitude (top) and power production (bottom) according to policy before (left), after half of training (center) and after full training (right) of agent H. Figure adapted from paper IV.

The policy resulting from training can be seen in Figure 3.15. With further training, the strategy reduces to setting constant values for the pitching amplitude of all three turbines. The power production by turbine 1 is slightly below that of the reference simulation, while power production by turbine two has increased significantly. However, this improvement can not be repeated at turbine 3. In total, the turbines

produce 6.8 % more power compared to the reference, mainly due to an increase in power production of 43 % by turbine 2, while power production by turbine 1 and 3 has decreased by 5 % and 8.8 %, respectively. As we show in our later work in paper V, the increase in power production observed for turbine 2 could not be repeated by turbine 3 because the relative angles of the helices lead to the destruction of the helical wake of the second turbine. This will be explored in more detail in section 3.5.

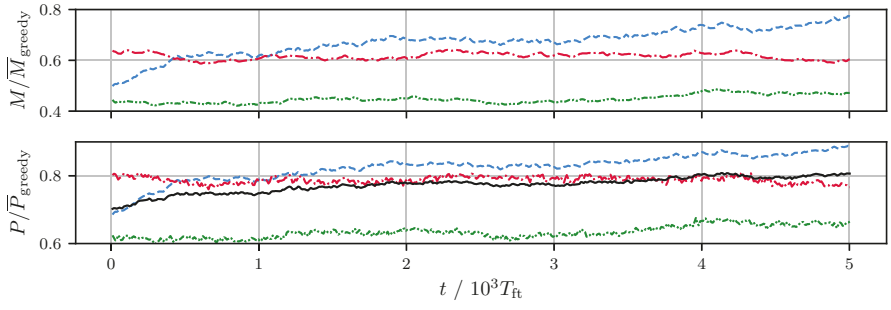

Figure 3.16. — Total, – – Turbine 1, –·– Turbine 2, ···· Turbine 3. Rolling average of generator torque (top) and power production (bottom) of the wind farm throughout the training of agent M-long. Respective powers normalized by mean per turbine of the reference simulation. Figure adapted from paper IV.

The evolution of agent M-long, the agent controlling the generator torque and emphasizing long-term returns, during training is shown in Figure 3.16. Training is carried out for 5000 flow-through times, equivalent to almost 7×10^6 s. Only the torque exerted at the first turbine increases, while it remains close to its initial values for the other two turbines. Power production by the first turbine also improves, while it is constant for turbine 2 and a small increase can be observed for turbine 3. Total power production remains at 80 % of the reference case. We therefore have to conclude that the training was not successful and the agent could not discover an equivalent or better strategy than greedy operation, which we attribute to the definition of the reward. The first turbine produces the most power, therefore increasing power production by the first turbine yields the most increase in reward, which is what PPO optimizes the policy for. The course of training agent M-short is shown in Figure 3.17. The exerted torque increases continuously, resulting in a continuous increase in power production by all three turbines. However, the total power production plateaus at the same level as the reference simulation. Furthermore, the first turbine is optimized first, then the second and third turbine improves last. Again, this is caused by the respective absolute contributions to the total power production. The evolution of the policy is presented in Figure 3.18. Before train-

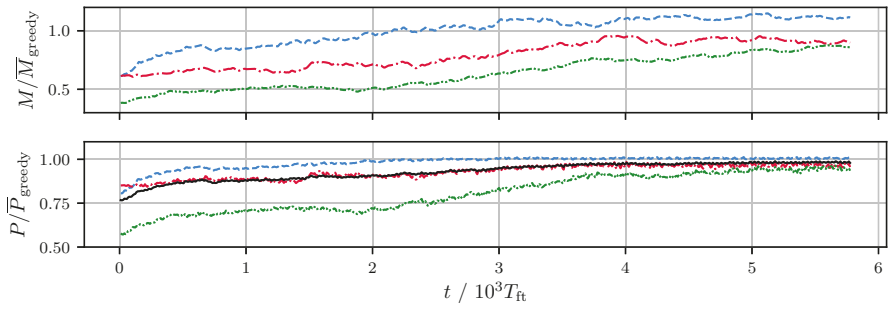

Figure 3.17. — Total, - - Turbine 1, -·- Turbine 2, -··- Turbine 3. Rolling average of generator torque (top) and power production (bottom) of the wind farm throughout the training of agent M-short. Respective powers normalized by mean per turbine of the reference simulation. Figure adapted from paper IV.

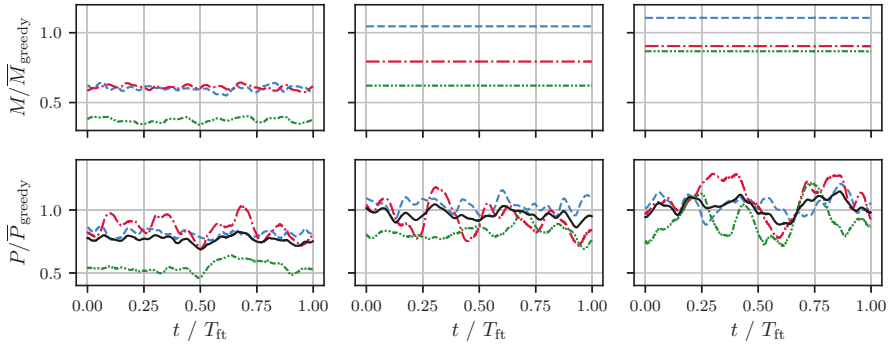

Figure 3.18. — Total, - - Turbine 1, -·- Turbine 2, -··- Turbine 3. Amplitude (top) and power production (bottom) according to policy before (left), after half of training (center) and after full training (right) of agent M-short. Figure adapted from paper IV.

ing the torque fluctuates, while after half of training it is constant for all three turbines. After the second half of training, the torque exerted at the first turbine has not changed much while the torque exerted on turbines two and three still increased. This reiterates that the first turbine is optimized first, then the second and so on. The agent develops a strategy similar to greedy control and finds similar values for the torque. However, the question remains, why the agent is not able to develop a strategy that exceeds the greedy control paradigm.

Looking back at the setup of the problem, multiple possible reasons emerge. One potential reason is the definition of the return. Recall, that the return is the discounted sum of rewards, see (2.16). In this case, the reward was defined as the total power production. The return at any

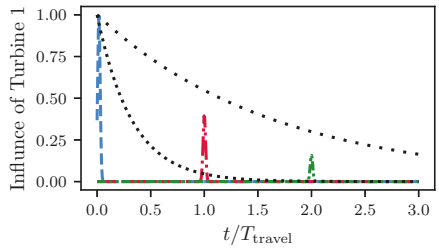

Figure 3.19. Turbine 1, Turbine 2, Turbine 3, $\gamma = 0.95$, $\gamma = 0.99$. Illustration of the influence of an action applied at turbine 1 compared to discounting factor with different discount rates. Figure adapted from paper IV.

timestep includes the influence of the current action on the respective turbine but also the influence of actions from timesteps in the past at the upstream turbines. The influence of an action at turbine 1 requires a time, T_{travel}, to reach the next downstream turbine. Therefore a single action influences the reward multiple times. In between, however, there is no influence. This is illustrated in Figure 3.19. If the return includes t and $t + T_{\text{travel}}$, as was the case for M-long, it must also include a lot of noise from the timesteps in between, therefore training is very slow. On the other hand, if the influence of the second turbine is disregarded, the policy must recede to a greedy policy. For future research, we therefore proposed to alter the definition of return to

$$R_t^1 = P_t^1 + P_{t+T_{\text{travel}}}^2 + P_{t+2T_{\text{travel}}}^3, \tag{3.1}$$

where the superscript indicates the turbine number.

While we ultimately were not able to discover a new control paradigm, we were able to show that it is possible to use reinforcement learning even in a high-fidelity LES environment. However, the computational cost of training is extremely high, even with LBM. With more computing power, multiple instances of the environment can be run simultaneously, as was done in Rabault and Kuhnle, 2019, and also to some extent for this work. The choice of letting the agent control the generator torque must also be reconsidered, as an increasing body of research points to the limited benefit achievable with static axial induction control, see Kheirabadi and Nagamune, 2019 for a review of wind farm control approaches. Recent work on yaw control with reinforcement learning by Liew et al., 2023 was shown to be successful, albeit in lower fidelity environments. Dynamic induction control achieved by actuation of collective pitch or individual pitch, i.e. the pulse and the helix approach were also very successful.

3.5 Characterizing the Helix Approach

With larger wind farms and larger turbines, wind farm control becomes a promising way to increase wind farm efficiency, ultimately reducing cost of electricity and minimizing detrimental environmental effects. As described in section 1.2, the two main mechanisms of wind farm control are axial induction control, aiming at reducing the wake effect, and wake steering, typically by yaw-misalignment. Recently, dynamic approaches, such as the helix approach by Frederik et al., 2020b, have been shown very promising results. Yet, in our previous study, we found that we could not replicate the increase in power production observed at the second turbine at the third turbine. Furthermore, the mechanism forming the helical wake and the effect the helical wake has were not yet examined. Therefore, paper V contains various theoretical considerations, on the formation and propagation of the helical wake, and numerical studies on the influence of tip-speed ratio, turbulence intensity and turbulent length scale on the wake. Furthermore, the interaction of multiple helices is examined theoretically and numerically.

We conducted 8 simulations of a single NREL 5MW turbine with the LBM solver `elbe`. The baseline case has an inflow velocity of 9 m/s, a turbulence intensity of 5%, and a turbulent length scale of 40 m. To perform a parameter study, either the inflow velocity is increased to 13 m/s to achieve a lower tip-speed ratio, the turbulence intensity is increased to 10 % or the inflow turbulence has a larger turbulent length scale of 120 m. Per case, one simulation with and without the helix approach is conducted. For further details on the cases and the control, I refer to the original publication. The names of the cases consist of a first letter indicating the method, ('G' for greedy and 'H' for helix), the number of turbines, which parameter is compared ('V' for velocity, 'I' for turbulence intensity or 'L' for turbulence length scale) and the value of that parameter, e.g. the case with helix control and the higher inflow velocity is called H1V13.

The turbulent inflow is constructed by superimposing a constant uniform inflow velocity with Mann type synthetic turbulence. The turbine has a rotor diameter of $D = 126$ m and the outer domain has a size of $8D \times 8D \times 15D$ and a grid spacing of $\Delta x = \frac{D}{16}$. Around the turbine a refined layer is placed with a grid spacing of $\frac{D}{32}$. We also conduct 5 simulations with three NREL 5MW turbines in a row in a larger domain. The inflow conditions are the same as the baseline case. One simulation, serving as reference, is conducted without the helix, and the other four are conducted with helices but varying phase difference. The naming scheme is the same as for the single turbine cases. The parameter is abbreviated as 'PS' for phase shift.

In the helix approach, the turbine blades are pitched sinusoidally

$$\Theta_b = A \sin(\Psi + \omega_e t), \tag{3.2}$$

, where Θ_b is the pitch of blade b, A is the pitch amplitude, ω_e is the excitation frequency of the helix and t is time. The azimuthal angle Ψ is zero when the blade is aligned with the z axis. The excitation frequency is non-dimensionalized as the Strouhal number

$$\mathrm{St} = \frac{2\pi\omega_e D}{V_0}, \tag{3.3}$$

where V_0 is the inflow velocity. A Strouhal number of 0.25 is used, which was found to be the optimal value by Frederik et al., 2020b. From these definitions, we can show that a radially rotating force F is exerted onto the flow, which is proportional to A and points in the direction $\phi = \pi - \omega_e t$. A detailed derivation can be found in appendix A of paper V.

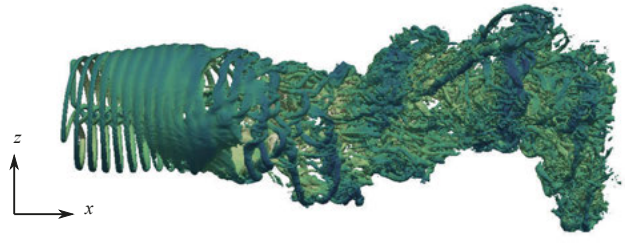

Figure 3.20. Vorticity contour in the wake of a turbine operated with the helix approach. Figure adapted from paper V.

To examine the effect of the helix approach on the tip-vortices in the wake, one more simulation of a single NREL 5MW turbine is performed with a higher resolution with a grid spacing of $\frac{D}{64}$ around the turbine. The vorticity contour in a snapshot of that simulation is shown in Figure 3.20. The changing pitch affects the size of the vortex core and leads to an early and uneven merging of vortices. This preempts the typical breakdown mechanism of tip-vortices. Since the presence of tip-vortices shields the wake from turbulent mixing, an earlier breakdown enhances the turbulent wake mixing, thus decreasing the wake effect. The distinct helical shape of the wake is also visible. The helical shape is not only due to the deformation of the wake but is actually due to a deflection of the whole wake.

To further identify the mechanisms at play in the helical wake, a more rigorous definition of the terms describing the wake is needed. We differentiate three phenomena: deformation, deflection and meandering of the wake. Deformation is a change of wake shape without translation of the

wake center. Deflection is a translation of the average wake center position due to forces exerted by the turbine. Crucially, we have to extend this definition to include the average wake center position in a rotating frame of reference that rotates with the excitation frequency of the helix. Finally, meandering describes the stochastic fluctuation of the wake center position due to large-scale turbulent structures in the incoming flow. Instantaneous flow fields were recorded at various downstream distances in the single turbine cases and identified the wake center position by finding the point where the velocity deficit is the best fit to a double Gaussian, see Quon et al., 2020 for a discussion on wake center identification methods. The wake center distance d follows a Rice-Distribution,

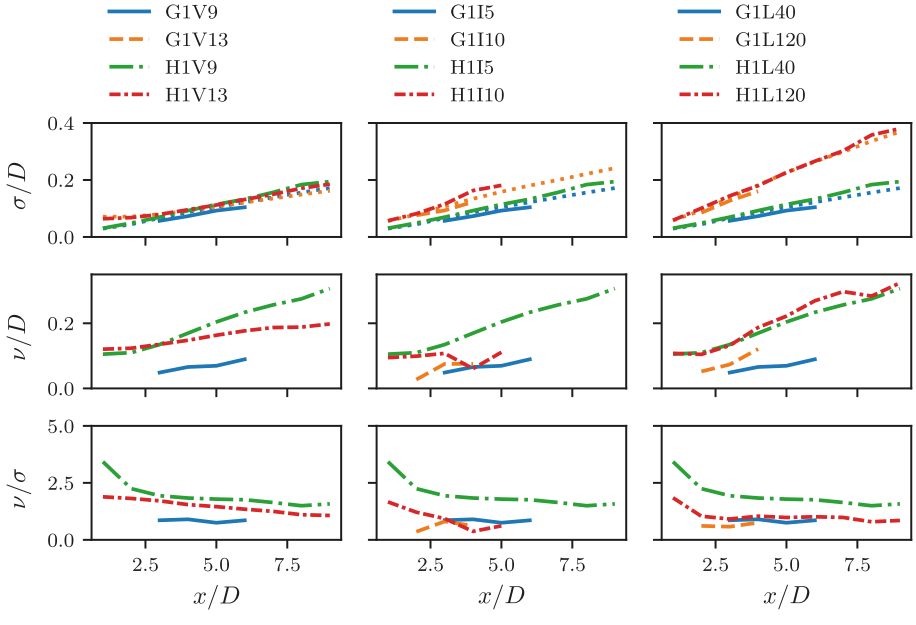

Figure 3.21. Parameters of the Rice distribution. (top) Meandering parameter, (center) deflection parameter, (bottom) ratio of deflection over meandering parameter. Dotted lines are results from fitting the meandering parameter for greedy cases based on a Rayleigh distribution. Figure adapted from paper V.

discussed by Rice, 1944, 1945, which has two parameters, ν and σ, where ν can be identified with the deflection and σ with the meandering of the wake. The result of fitting the parameters to the calculated wake center distances is shown in Figure 3.21. The meandering is not affected by the presence of the helix. However, the deflection parameter in the helix cases grows with downstream distance. For $\nu \to 0$, the algorithm for the parameter identification becomes unstable. Since the algorithm did not converge for the greedy cases, we can assume that the deflection was close to zero. The dotted lines show a fit of σ in the greedy cases according

to a Rayleigh distribution. The Rayleigh distribution is a degenerated Rice distribution with $\nu = 0$. The results for σ are nearly identical to the results from fitting the Rice parameters. Comparing the cases with differing inflow velocity in the left column shows that the meandering is not affected by an increase in inflow velocity, while the strength of the deflection decreases. In the center column, the algorithm also did not converge for the cases with high turbulence intensity, since the high turbulence intensity makes the wake center identification difficult. The larger turbulence length scale amplifies wake meandering, as has been observed in many previous studies, such as Larsen et al., 2008, while not greatly affecting the wake deflection. In summary, the helix approach causes wake deflection, while not affecting the meandering of the wake.

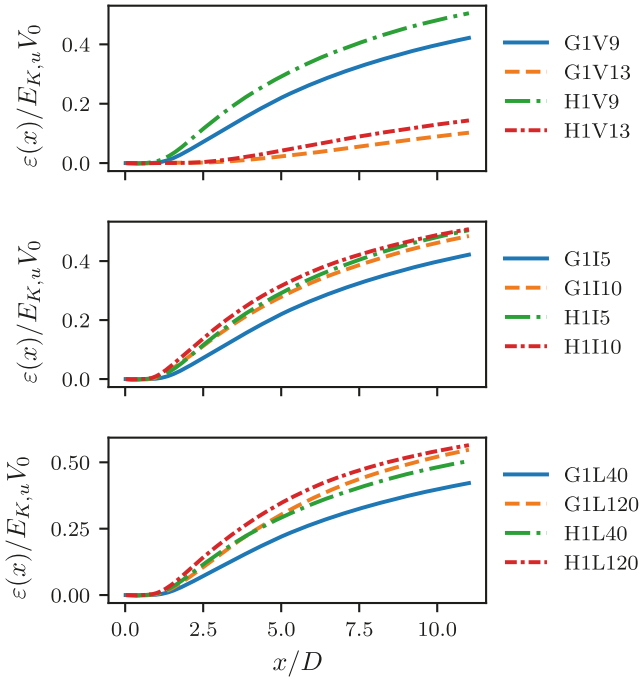

Figure 3.22. Cumulative entrainment of kinetic energy into the cylinder behind the rotor swept area. Comparing (top) inflow velocity, (center) turbulence intensity and (bottom) turbulent length scale. Figure adapted from paper V.

The earlier tip-vortex breakdown and consequently faster wake mixing as well as the wake deflection lead to more energy in the area behind the turbine. Figure 3.22 shows the cumulative entrainment of kinetic energy into the cylinder behind the turbine rotor, i.e. the flux of energy into the rotor swept area of a hypothetical downstream turbine. The cumulative entrainment is denoted as $\epsilon(x)$, while the kinetic energy of the undisturbed flow is denoted as $E_{K,u}$. Entrainment increases in all cases

due to the helix approach, however, that increase is most pronounced in the baseline case. In general, less energy is entrained at the higher inflow velocity, however, the turbine also extracts less energy, relatively speaking. At higher turbulence intensities and larger turbulent length scales, entrainment is higher even without the helix, yet the application of the helix increases entrainment even further.

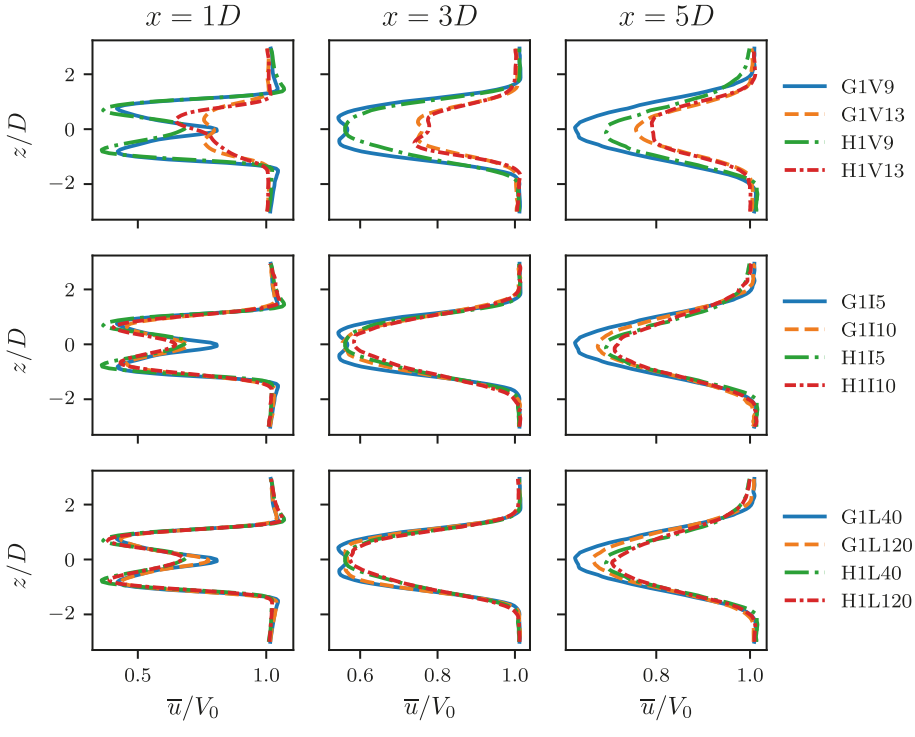

Figure 3.23. Profiles of average inflow velocity at different downstream distances, average computed in the meandering frame of reference at three different downstream distances. Comparing (top) inflow velocity, (center) turbulence intensity and (bottom) turbulent length scale. Figure adapted from paper V.

The increased entrainment results in higher inflow velocity, as can be seen in Figure 3.23. As Figure 3.22 has already shown, the helix approach is most effective at low velocity, low turbulence intensity, and low turbulence length scale. Nevertheless, an increase in inflow velocity is present in all cases. Note that the velocity is averaged in a meandering frame of reference. The frame of reference is rotated with the helix and shifted so that the wake center is in the center of the plane, pronouncing the asymmetric wake shape. In the original publication the wake shape is examined even further. The wake consists of two vortices, similar to findings in yawed wakes, however, the vortices are of different strengths.

We also develop a kinematic description of the helical wake based on a helix transport velocity

$$\overline{u}_{\text{helix}}(x) = \frac{x - x_0}{t - t_0}, \tag{3.4}$$

where x_0 is the location of the turbine and t_0 is the time a fluid particle passes the rotor-swept area. Consequently, the angle of the helix is

$$\varphi(x,t) - \varphi(x_0,t) = 2\pi \mathrm{St} \frac{x - x_0}{D} \frac{V_0}{\overline{u}_{helix}} = \Delta\varphi(x). \tag{3.5}$$

Measurements of $\Delta\varphi(x)$ and $\overline{u}_{helix}$ in the single turbine simulations show that they are not affected by changes in the inflow turbulence, only at the higher inflow velocity, which has to be expected. $\Delta\varphi(x)$ increases almost linearly with downstream distance.

The previous observations allow us to develop a description of two turbines with helical wakes. First, it has to be established that both turbines have to operate the helix approach at the same ω_e to achieve a constant phase difference between the helices. Furthermore, it is noted that the phase shift can be set explicitly by extending (3.2) to

$$\Theta = A \sin\left(\Psi_b + \omega_e t - \left(2\pi \mathrm{St} \frac{S_x}{D} \frac{V_0}{\overline{u}_{helix}} + \Phi\right)\right), \tag{3.6}$$

where S_x is the distance between the turbines and Φ is the desired phase shift.

We then turn to examining the results from the three turbine simulations. First, the accuracy of the a-priori-determined phase shifts is studied. Due to induction, the a-priori-determined phase shifts are off by 19° to 36°.

Examination of the angle of attack the blades of the second turbine shows that the optimal angle of attack along the whole blade is achieved at a phase shift of 0°, while a phase shift of 180° leads to large variations in the angle of attack and thus suboptimal lift and drag forces on the blade.

The wake center distribution is also investigated. The distribution in the wake of the first turbine is identical in all for cases with the helix, while the distribution in the wake of the second turbine changes significantly depending on the phase shift. In the case of a phase shift of 90°, no deflection can be observed, while the deflection has increased in strength in the case with a phase shift of 270°.

The different strengths in wake deflection also result in large differences in the kinetic energy in the flow, which is shown in Figure 3.24, where $\overline{E}_{K,R}$ is the mean flow kinetic energy in the rotor disc. There is no difference between the helix cases in the wake of the first turbine, while

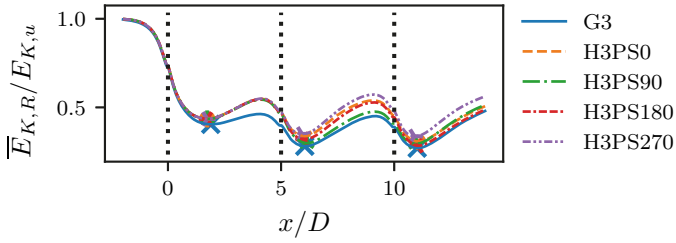

Figure 3.24. Mean kinetic energy in the rotor disk of all three turbine cases. Figure adapted from paper V.

there is significantly less energy in the wake in the reference case. In the wake of the second turbine, the maximum kinetic energy is even higher than the maximum in the wake of the first turbine in case H3PS270, while the energy in case H3PS90 is only slightly higher than in the reference case. The wake of the third turbine, where the phase shift was not optimized, follows a similar trend. This indicates, that the phase shift has great influence on the achievable power gains with the helix approach and that the phase shift must be optimized for each turbine. However, since turbulence conditions have little influence on the helix transport velocity, the phase shifts should also be stable under changing turbulence conditions. Furthermore, as long as the turbine operates at a constant tip-speed ratio, the ratio of $\overline{u}_{helix}/V_0$ should also stay constant, thus the phase shifts also stay constant with changing inflow velocity.

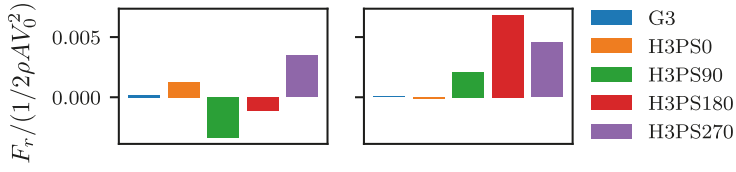

Figure 3.25. Radial force exerted by the second turbine, split into component aligned with the incoming helix (left) and perpendicular component (right).

In order to find the reason that H3PS270 has the highest kinetic energy, Figure 3.25 shows the radial force exerted by the second turbine. In case H3PS90, a strong force in the opposite direction of the deflection of the incoming helix destroys the helix, while the incoming helix is amplified by a strong force in the same direction in case H3PS270. The strongest force is found in case H3PS180, however, it acts at a right angle to the incoming helic and is thus less effective.

Finally, the power and thrust of the individual turbines and the total power are presented in Figure 3.26. Applying the helix approach leads to a small decrease in power production by the first turbine compared to the

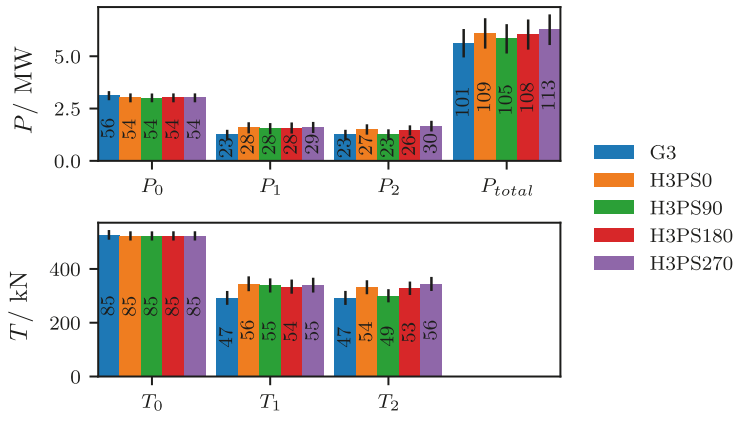

Figure 3.26. Average power (top) and thrust (bottom) of the turbines of all three turbine cases. Black lines indicate one standard deviation above and below average. Numbers power and thrust coefficients. Figure adapted from paper V.

reference case. However, the second turbine produces five to six percent of available kinetic energy more. Due to the increased inflow velocity the thrust is increased as well. The third turbine also produces more power in all cases except H3PS90, but case H3PS270 has the highest power production, as was indicated by all previous results. In total, the power production increases by 12% in case H3PS270 and even in case H3PS90 the total power is increased by 4%.

We examined many aspects of the helix, yet, there are important effects that were not considered. First, our simulations were conducted under highly idealized conditions in order to isolate effects of the helix. However, the presence of the ground and shear will influence the effectiveness of the approach, especially the wake deflection. Furthermore, the actuation of the pitch increases the fatigue loads of the turbine. Recent work by Taschner et al., 2023 elucidates both the effect of the ground on wake center deflection and examines the relation of pitch angle and damage equivalent loads. Further research should also explore ways of finding the optimal phase shift. Additionally, research on the optimal value for the Strouhal number is still inconclusive, Frederik et al., 2020b reported an optimal value of 0.25, while Muscari et al., 2022 finds 0.4 to be optimal. Physical reasons, related to the stability of the wake for example, have not yet been published. Overall, our results help understand the mechanics of the helix approach, and provide tools to describe the helical wake more thoroughly. We show that the power increase stems from enhanced wake mixing but also wake deflection. We also present an extension of the approach to more than one turbine considering the phase shift between the helices.

4. Conclusion and Outlook

The recent introduction of LBM to the field of wind energy enables many potential new use cases. The goal of this thesis is to develop LBM into a versatile tool for the field of wind energy, demonstrate the LBM's accuracy and efficiency, and how it opens up new applications for LES. I showed how the industrial application of LES becomes feasible with current methods, demonstrated that the method is accurate while being multiple orders of magnitude more efficient than traditional LES solvers and how the combination of LBM with different machine learning methods can lead to even faster modeling tools or enables new data-driven approaches. The method can also be used to aid in the understanding of physical mechanisms, which resulted in showing how the helix approach can be used to increase wind farm efficiency by up to 12% in a specific case.

Industrial experts saw many benefits in the use of LES for wind farm simulation, especially with phenomena associated with large offshore wind farms and onshore wind farms in complex terrain. Yet, the computational cost and long run times hinder the adoption of LES. Other factors include a lack of trained staff and validation. In contrast, we showed that with the current state-of-the-art LBM simulation, the runtime demands by about half of the respondents could be met for a case similar to the indicated typical use case.

We validated our methodology against data from the SWiFT benchmark and other LES simulations of the same case. We simulated a single Vestas V27 turbine using an LBM solver coupled to a wind turbine simulation tool that uses the actuator line method and a smearing correction. We showed that the methodology accurately predicts the turbine response and several wake quantities, including velocity deficit, wake strength, recovery, and meandering. A comparison of the energy to solution showed that our method is at least one order of magnitude, in most cases more than two orders of magnitude, more efficient than other LES methods, while being one order of magnitude more costly than simulations using RANS.

We also used LBM to create a training dataset of the flow around a single turbine. The data was used to train a convolutional neural network, that, based on inflow velocity, turbulence intensity, and some turbine data, predicts the average flow field and turbulence intensity

downstream of the turbine. The trained network showed excellent agreement when tested with never-before-seen data while executing in a few microseconds on a consumer-grade GPU.

Another machine learning approach that has been successful at finding control strategies is reinforcement learning. We trained three different agents in an environment of three turbines subjected to turbulent inflow. One of the agents controlled the pitch amplitude of the turbines applying the helix approach, while the two others controlled the generator torque of the turbines, one with emphasis on long-time reward, while the other emphasized short-term rewards. The algorithm successfully improved the pitch angles of the first agent. The second agent, however, did not converge to a strategy that reached the performance of a greedily operated farm and the third agent could also not find a strategy exceeding greedy control and rediscovered a state close to greedy control. One possible explanation is the definition of the training objective, which was defined without taking the time delay between the turbines into account.

Finally, we further examined the helix approach used in the previous study. We highlighted the mechanism causing the helical wake and showed that the helix approach leads to earlier tip-vortex breakdown, consequently enhancing wake mixing, but also causes a deflection of the wake. Both effects together lead to improved entrainment of kinetic energy into the wake. We provided a kinematic description of the wake and showed that the helix transport velocity is not affected by changes in the turbulence characteristics. We also noted some considerations on the interaction of multiple turbines using the helix approach and found an expression to explicitly set the phase difference between two helical wakes. From there we conducted a parameter study on the effect of the phase difference, showing that a phase difference of 270° yields the strongest deflection of the second wake and consequently the most kinetic energy in the wake of the second and third turbine. We explained the varying strength of the helix with the strength and relative direction of the radial force at the second turbine, thus a combination of optimal angle of attack and phase of the helix. We found that the helix results in an increase in power production in all cases, although the gains are limited to the second turbine in the least favorable case. The optimal case achieved an improvement of 12% compared with the reference case where all turbines operate in greedy mode.

While this work covers many aspects important in further developing LBM for wind energy research there remains a plethora of further research topics. As paper I shows, complex terrain and accurate simulation of large offshore wind farms are of high interest. Yet, crucial models to accomplish such simulations are missing as of now, namely an extension of the iMEM boundary condition to complex geometries and thermal stratification. Furthermore, validation with larger, flexible

turbines is missing, however, there exists no public validation data for such a case. This will also require the incorporation of further models to account for the flexibility of the blade. Beyond traditional methods of validation, data-driven approaches such as uncertainty quantification might also become feasible with LBM, further increasing trust in LBM and LES in general. Further building up training datasets, including waked turbines, for example, will help train more machine learning models, aiding in applications where engineering models are used today. The reduced computational demand also allows to conduct larger suits of simulations, to conduct parameter studies, for example, improving our understanding of the physics of wind farm flows.

5. Samanfattning

Många nya utmaningar inom vindenergi kräver stora datorberäkningar med den s.k. large eddy-metoden för att nå den noggrannhet som krävs när flöden i vindkraftparker beräknas. Den enorma beräkningskostnaden hindrar dock att den används inom forskning och industri. Lattice Boltzmann-metoden är den mest lovande kandidaten hittills för att uppnå en så hög nivå av noggrannhet och samtidigt minska beräkningskostnaderna, jämfört med traditionella metoder. I denna avhandling presenterar jag en vidareutveckling av Lattice Boltzmann-metoden för vindenergiområdet och olika applikationer såsom industriell användning, generering av träningsdata för maskininlärning och analys av kontrollrutiner för parkstyrning.

För att utvärdera kraven från olika industriella intressenter genomfördes en undersökning bland branschexperter om användningen av large eddy-simuleringar och visade att kraven på beräkningskostnad, som angavs av många deltagare, kan uppfyllas genom användning av dagens utvecklingsnivå av Lattice Boltzmann-metoden. I en valideringsstudie visade sig Lattice Boltzmann-metoden att vara lika exakt som traditionella Navier-Stokes-lösare samtidigt som de minskar beräkningskostnaderna med en till fler storleksordningar. Konvolutionella neurala nätverk tränades att förutsäga genomsnittliga strömningshastigheter i vaken bakom en turbin. Förutsägelserna uppvisade mycket hög noggrannhet vid en beräkningstid motsvarande ingenjörsmetoder. Lattice Boltzmann-metoden möjliggör därför generering av ett större antal träningsuppsättningar till en genomförbar beräkningskostnad. Ett ”proof of concept” tillhandahålls för användning av s.k. reinforcement learning för att upptäcka nya kontrollmekanismer för parkstyrning. En omfattande analys av helixmetoden presenteras där dess fysiska mekanismer belyses och en grundlig parameterstudie av en enskild turbins vak genomförts. Vi studerar också interaktionen mellan flera spiralformade vakar som ger ett sätt att utöka tillvägagångssättet från ett par turbiner till hela vindkraftsparker.

Acknowledgments

I would like to thank my current and past supervisors Stefan Ivanell, Joakim Widén, Karl Nilsson and Hugo Olivares-Espinosa for their encouragement, support and guidance. I also thank Henrik Asmuth, who has been a wonderful colleague and friend and has paved so many ways for me.

I am also grateful for all my other colleagues at the Wind Energy Section for their support and their understanding. My gratitude goes to everyone at the Graduate School for Sustainability Studies, especially Jenny Helin, for creating a space filled with generosity and kindness.

Bibliography

Asmuth, Henrik, Christian F. Janßen, Hugo Olivares-Espinosa, and Stefan Ivanell (2021). "Wall-Modeled Lattice Boltzmann Large-Eddy Simulation of Neutral Atmospheric Boundary Layers". In: *Physics of Fluids* 33.10, p. 105111.

Asmuth, Henrik, Gonzalo P. Navarro Diaz, Helge Aagaard Madsen, Emmanuel Branlard, Alexander R. Meyer Forsting, Karl Nilsson, Jason Jonkman, and Stefan Ivanell (2022). "Wind Turbine Response in Waked Inflow: A Modelling Benchmark against Full-Scale Measurements". In: *Renewable Energy* 191, pp. 868–887.

Asmuth, Henrik, Hugo Olivares-Espinosa, and Stefan Ivanell (2020). "Actuator Line Simulations of Wind Turbine Wakes Using the Lattice Boltzmann Method". In: *Wind Energy Science* 5.2, pp. 623–645.

Asmuth, Henrik, Hugo Olivares-Espinosa, Karl Nilsson, and Stefan Ivanell (2019). "The Actuator Line Model in Lattice Boltzmann Frameworks: Numerical Sensitivity and Computational Performance". In: *Journal of Physics: Conference Series* 1256, p. 012022.

Bak, Christian, Robert Bitsche, Anders Yde, Taeseong Kim, Morten Hartvig Hansen, Frederik Zahle, Mac Gaunaa, José Pedro Albergaria Amaral Blasques, Mads Døssing, Jens-Jakob Wedel Heinen, et al. (2012). "Light Rotor: The 10-MW Reference Wind Turbine". In: *EWEA 2012-European Wind Energy Conference & Exhibition*. European Wind Energy Association (EWEA).

Berg, Daniel van den, Delphine de Tavernier, and Jan-Willem van Wingerden (2022). "Using The Helix Mixing Approach On Floating Offshore Wind Turbines". In: *Journal of Physics: Conference Series* 2265.4, p. 042011.

Bhatnagar, P. L., E. P. Gross, and M. Krook (1954). "A Model for Collision Processes in Gases. I. Small Amplitude Processes in Charged and Neutral One-Component Systems". In: *Physical Review* 94.3, pp. 511–525.

Bouzidi, M'hamed, Mouaouia Firdaouss, and Pierre Lallemand (2001). "Momentum Transfer of a Boltzmann-lattice Fluid with Boundaries". In: *Physics of Fluids* 13.11, pp. 3452–3459.

Bove, M, B López, J Toutouh, S Nesmachnow, and M Draper (2023). "Estimation of Wind Turbine Wakes with Generative-Adversarial Networks". In: *Journal of Physics: Conference Series* 2505.1, p. 012053.

Breton, S.-P., J. Sumner, J. N. Sørensen, K. S. Hansen, S. Sarmast, and S. Ivanell (2017). "A Survey of Modelling Methods for High-Fidelity Wind Farm Simulations Using Large Eddy Simulation". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 375.2091, p. 20160097.

Coreixas, Christophe, Bastien Chopard, and Jonas Latt (2019). "Comprehensive Comparison of Collision Models in the Lattice Boltzmann Framework: Theoretical Investigations". In: *Physical Review E* 100.3, p. 033305.

Coreixas, Christophe, Gauthier Wissocq, Bastien Chopard, and Jonas Latt (2020). "Impact of Collision Models on the Physical Properties and the Stability of Lattice Boltzmann Methods". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 378.2175, p. 20190397.

Dağ, Kaya Onur and Jens Nørkær Sørensen (2020). "A New Tip Correction for Actuator Line Computations". In: *Wind Energy* 23.2, pp. 148–160.

Deardorff, J. W. (1972). "Numerical Investigation of Neutral and Unstable Planetary Boundary Layers". In: *Journal of the Atmospheric Sciences* 29.1, pp. 91–115.

Demuth, Howard B., Mark H. Beale, Orlando De Jess, and Martin T. Hagan (2014). *Neural Network Design*. 2nd. Stillwater, OK, USA: Martin Hagan.

Doubrawa, Paula, Mithu Debnath, Patrick J. Moriarty, Emmanuel Branlard, Thomas G. Herges, Dave Maniaci, and Jonathan Naughton (2019). "Benchmarks for Model Validation Based on LiDAR Wake Measurements". In: *Journal of Physics: Conference Series* 1256.1, p. 012024.

Doubrawa, Paula, Eliot W. Quon, Luis A. Martinez-Tossas, Kelsey Shaler, Mithu Debnath, Nicholas Hamilton, Thomas G. Herges, Dave Maniaci, Christopher L. Kelley, Alan S. Hsieh, Myra L. Blaylock, Paul van der Laan, Søren Juhl Andersen, Sonja Krueger, Marie Cathelain, Wolfgang Schlez, Jason Jonkman, Emmanuel Branlard, Gerald Steinfeld, Sascha Schmidt, Frédéric Blondel, Laura J. Lukassen, and Patrick Moriarty (2020). "Multimodel Validation of Single Wakes in Neutral and Stratified Atmospheric Conditions". In: *Wind Energy* 23.11, pp. 2027–2055.

Frederik, Joeri Alexis, Bart M. Doekemeijer, Sebastiaan P. Mulders, and Jan-Willem van Wingerden (2020a). "On Wind Farm Wake Mixing Strategies Using Dynamic Individual Pitch Control". In: *Journal of Physics: Conference Series* 1618.2, p. 022050.

Frederik, Joeri Alexis, Bart M. Doekemeijer, Sebastiaan P. Mulders, and Jan-Willem van Wingerden (2020b). "The Helix Approach: Using Dynamic Individual Pitch Control to Enhance Wake Mixing in Wind Farms". In: *Wind Energy* 23.8, pp. 1739–1751.

Fukushima, Kunihiko (1980). "Neocognitron: A Self-Organizing Neural Network Model for a Mechanism of Pattern Recognition Unaffected by Shift in Position". In: *Biological Cybernetics* 36.4, pp. 193–202.

Gaudet, Brian, Richard Linares, and Roberto Furfaro (2020). "Deep Reinforcement Learning for Six Degree-of-Freedom Planetary Landing". In: *Advances in Space Research* 65.7, pp. 1723–1741.

Geier, Martin, Andrea Pasquali, and Martin Schönherr (2017a). "Parametrization of the Cumulant Lattice Boltzmann Method for Fourth Order Accurate Diffusion Part I: Derivation and Validation". In: *Journal of Computational Physics* 348, pp. 862–888.

Geier, Martin, Andrea Pasquali, and Martin Schönherr (2017b). "Parametrization of the Cumulant Lattice Boltzmann Method for Fourth Order Accurate Diffusion Part II: Application to Flow around a Sphere at Drag Crisis". In: *Journal of Computational Physics* 348, pp. 889–898.

Geier, Martin, Martin Schönherr, Andrea Pasquali, and Manfred Krafczyk (2015). "The Cumulant Lattice Boltzmann Equation in Three Dimensions: Theory and Validation". In: *Computers & Mathematics with Applications* 70.4, pp. 507–547.

Goit, Jay P. and Johan Meyers (2015). "Optimal Control of Energy Extraction in Wind-Farm Boundary Layers". In: *Journal of Fluid Mechanics* 768, pp. 5–50.

Goit, Jay P., Wim Munters, and Johan Meyers (2016). "Optimal Coordinated Control of Power Extraction in LES of a Wind Farm with Entrance Effects". In: *Energies* 9.1, p. 29.

Goodfellow, Ian, Yoshua Bengio, and Aaron Courville (2016). *Deep Learning*. Adaptive Computation and Machine Learning. Cambridge, Massachusetts: The MIT Press.

Hansen, Martin O. L. (2008). *Aerodynamics of Wind Turbines*. 2nd ed. London ; Sterling, VA: Earthscan.

Hochreiter, Sepp and Jürgen Schmidhuber (1997). "Long Short-Term Memory". In: *Neural Computation* 9.8, pp. 1735–1780.

IEA (2022). *World Energy Outlook 2022*. Tech. rep.

Ivanell, Stefan (2009). "Numerical Computations of Wind Turbine Wakes". In:

Jézéquel, E, M Cathelain, V Masson, and F Blondel (2021). "Validation of Wind Turbine Wakes Modelled by the Meso-NH LES Solver under Different Cases of Stability". In: *Journal of Physics: Conference Series* 1934.1, p. 012003.

Jonkman, J., S. Butterfield, W. Musial, and G. Scott (2009). *Definition of a 5-MW Reference Wind Turbine for Offshore System Development*. Tech. rep. NREL/TP-500-38060. National Renewable Energy Laboratory.

Khanna, Samir and James G. Brasseur (1997). "Analysis of Monin–Obukhov Similarity from Large-Eddy Simulation". In: *Journal of Fluid Mechanics* 345, pp. 251–286.

Kheirabadi, Ali C. and Ryozo Nagamune (2019). "A Quantitative Review of Wind Farm Control with the Objective of Wind Farm Power Maximization". In: *Journal of Wind Engineering and Industrial Aerodynamics* 192, pp. 45–73.

Kingma, Diederik P. and Jimmy Ba (2017). "Adam: A Method for Stochastic Optimization". In: *arXiv:1412.6980 [cs]*.

Kleine, Vitor G., Ardeshir Hanifi, and Dan S. Henningson (2023). "Non-Iterative Vortex-Based Smearing Correction for the Actuator Line Method". In: *Journal of Fluid Mechanics* 961, A29.

Krüger, Timm, Halim Kusumaatmaja, Alexandr Kuzmin, Orest Shardt, Goncalo Silva, and Erlend Magnus Viggen (2017). *The Lattice Boltzmann Method: Principles and Practice*. 1st ed. Graduate Texts in Physics. Springer International Publishing.

Ladd, Anthony J. C. (1994). "Numerical Simulations of Particulate Suspensions via a Discretized Boltzmann Equation. Part 1. Theoretical Foundation". In: *Journal of Fluid Mechanics* 271, pp. 285–309.

Lallemand, Pierre and Li-Shi Luo (2003). "Lattice Boltzmann Method for Moving Boundaries". In: *Journal of Computational Physics* 184.2, pp. 406–421.

Larsen, Gunner C., Helge Aa. Madsen, Kenneth Thomsen, and Torben J. Larsen (2008). "Wake Meandering: A Pragmatic Approach". In: *Wind Energy* 11.4, pp. 377–395.

Latt, Jonas, Bastien Chopard, Orestis Malaspinas, Michel Deville, and Andreas Michler (2008). "Straight Velocity Boundaries in the Lattice Boltzmann Method". In: *Physical Review E* 77.5, p. 056703.

Lee, S., M. J. Churchfield, P. J. Moriarty, J. Jonkman, and J. Michalakes (2013). "A Numerical Study of Atmospheric and Wake Turbulence Impacts on Wind Turbine Fatigue Loadings". In: *Journal of Solar Energy Engineering* 135.3.

Li, Rui, Jincheng Zhang, and Xiaowei Zhao (2022). "Dynamic Wind Farm Wake Modeling Based on a Bilateral Convolutional Neural Network and High-Fidelity LES Data". In: *Energy* 258, p. 124845.

Liew, Jaime, Tuhfe Göçmen, Wai Hou Lio, and Gunner Chr. Larsen (2023). "Model-free Closed-loop Wind Farm Control Using Reinforcement Learning with Recursive Least Squares". In: *Wind Energy*, we.2852.

Lignarolo, Lorenzo E. M., Dhruv Mehta, Richard J. A. M. Stevens, Ali Emre Yilmaz, Gijs van Kuik, Søren J. Andersen, Charles Meneveau, Carlos J. Ferreira, Daniele Ragni, Johan Meyers, Gerard J. W. van Bussel, and Jessica Holierhoek (2016). "Validation of Four LES and a Vortex Model against Stereo-PIV Measurements in the near Wake

of an Actuator Disc and a Wind Turbine". In: *Renewable Energy* 94, pp. 510–523.
Löhner, Rainald (2019). "Towards Overcoming the LES Crisis". In: *International Journal of Computational Fluid Dynamics* 33.3, pp. 87–97.
Malaspinas, Orestis and Pierre Sagaut (2014). "Wall Model for Large-Eddy Simulation Based on the Lattice Boltzmann Method". In: *Journal of Computational Physics* 275, pp. 25–40.
Mann, Jakob (1998). "Wind Field Simulation". In: *Probabilistic Engineering Mechanics* 13.4, pp. 269–282.
Martínez-Tossas, Luis A., Matthew J. Churchfield, Ali Emre Yilmaz, Hamid Sarlak, Perry L. Johnson, Jens N. Sørensen, Johan Meyers, and Charles Meneveau (2018). "Comparison of Four Large-Eddy Simulation Research Codes and Effects of Model Coefficient and Inflow Turbulence in Actuator-Line-Based Wind Turbine Modeling". In: *Journal of Renewable and Sustainable Energy* 10.3, p. 033301.
Martínez-Tossas, Luis A. and Charles Meneveau (2019). "Filtered Lifting Line Theory and Application to the Actuator Line Model". In: *Journal of Fluid Mechanics* 863, pp. 269–292.
Mehta, D., A. H. van Zuijlen, B. Koren, J. G. Holierhoek, and H. Bijl (2014). "Large Eddy Simulation of Wind Farm Aerodynamics: A Review". In: *Journal of Wind Engineering and Industrial Aerodynamics* 133, pp. 1–17.
Meyer Forsting, Alexander R., Georg Raimund Pirrung, and Néstor Ramos-García (2019). "A Vortex-Based Tip/Smearing Correction for the Actuator Line". In: *Wind Energy Science* 4.2, pp. 369–383.
Meyer Forsting, Alexander R., Georg Raimund Pirrung, and Néstor Ramos-García (2020). "Brief Communication: A Fast Vortex-Based Smearing Correction for the Actuator Line". In: *Wind Energy Science* 5.1, pp. 349–353.
Mikkelsen, Robert F. (2003). "Actuator Disc Methods Applied to Wind Turbines". PhD thesis. Technical University Denmark.
Mikkelsen, Torben, N. Angelou, K. Hansen, M. Sjöholm, M. Harris, C. Slinger, P. Hadley, R. Scullion, G. Ellis, and G. Vives (2013). "A Spinner-integrated Wind Lidar for Enhanced Wind Turbine Control". In: *Wind Energy* 16.4, pp. 625–643.
Munters, Wim and Johan Meyers (2017). "An Optimal Control Framework for Dynamic Induction Control of Wind Farms and Their Interaction with the Atmospheric Boundary Layer". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 375.2091, p. 20160100.
Munters, Wim and Johan Meyers (2018). "Towards Practical Dynamic Induction Control of Wind Farms: Analysis of Optimally Controlled

Wind-Farm Boundary Layers and Sinusoidal Induction Control of First-Row Turbines". In: *Wind Energy Science* 3.1, pp. 409–425.

Muscari, C, P Schito, A Viré, A Zasso, D van der Hoek, and Jw van Wingerden (2022). "Physics Informed DMD for Periodic Dynamic Induction Control of Wind Farms". In: *Journal of Physics: Conference Series* 2265.2, p. 022057.

Park, Jinkyoo, Soon-Duck Kwon, and Kincho Law (2017). "A Data-Driven, Cooperative Approach for Wind Farm Control: A Wind Tunnel Experimentation". In: *Energies* 10.7, p. 852.

Porté-Agel, Fernando, Majid Bastankhah, and Sina Shamsoddin (2020). "Wind-Turbine and Wind-Farm Flows: A Review". In: *Boundary-Layer Meteorology* 174.1, pp. 1–59.

Quon, Eliot W., P. Doubrawa, and M. Debnath (2020). "Comparison of Rotor Wake Identification and Characterization Methods for the Analysis of Wake Dynamics and Evolution". In: *Journal of Physics: Conference Series* 1452.1, p. 012070.

Rabault, Jean, Miroslav Kuchta, Atle Jensen, Ulysse Réglade, and Nicolas Cerardi (2019). "Artificial Neural Networks Trained through Deep Reinforcement Learning Discover Control Strategies for Active Flow Control". In: *Journal of Fluid Mechanics* 865, pp. 281–302.

Rabault, Jean and Alexander Kuhnle (2019). "Accelerating Deep Reinforcement Learning Strategies of Flow Control through a Multi-Environment Approach". In: *Physics of Fluids* 31.9, p. 094105.

Rice, S. O. (1944). "Mathematical Analysis of Random Noise". In: *The Bell System Technical Journal* 23.3, pp. 282–332.

Rice, S. O. (1945). "Mathematical Analysis of Random Noise". In: *The Bell System Technical Journal* 24.1, pp. 46–156.

Rozema, Wybe, Hyun J. Bae, Parviz Moin, and Roel Verstappen (2015). "Minimum-Dissipation Models for Large-Eddy Simulation". In: *Physics of Fluids* 27.8, p. 085107.

Sagaut, Pierre (2010). "Toward Advanced Subgrid Models for Lattice-Boltzmann-based Large-eddy Simulation: Theoretical Formulations". In: *Computers & Mathematics with Applications*. Mesoscopic Methods in Engineering and Science 59.7, pp. 2194–2199.

Sanderse, B., S.p. van der Pijl, and B. Koren (2011). "Review of Computational Fluid Dynamics for Wind Turbine Wake Aerodynamics". In: *Wind Energy* 14.7, pp. 799–819.

Schulman, John, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov (2017). "Proximal Policy Optimization Algorithms". In: *arXiv:1707.06347 [cs]*.

Serrano González, Javier, Manuel Burgos Payán, Jesús Riquelme Santos, and Ángel Gaspar González Rodríguez (2015). "Maximizing the Overall Production of Wind Farms by Setting the Individual Operating Point of Wind Turbines". In: *Renewable Energy* 80, pp. 219–229.

Shaw, Roger H. and Ulrich Schumann (1992). "Large-Eddy Simulation of Turbulent Flow above and within a Forest". In: *Boundary-Layer Meteorology* 61.1, pp. 47–64.

Silver, David, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis (2018). "A General Reinforcement Learning Algorithm That Masters Chess, Shogi, and Go through Self-Play". In: *Science* 362.6419, pp. 1140–1144.

Singh, Narender, Jeroen D. M. De Kooning, and Lieven Vandevelde (2022). "Dynamic Wake Analysis of a Wind Turbine Providing Frequency Support Services". In: *IET Renewable Power Generation* 16.9, pp. 1853–1865.

Smagorinsky, J. (1963). "GENERAL CIRCULATION EXPERIMENTS WITH THE PRIMITIVE EQUATIONS: I. THE BASIC EXPERIMENT". In: *Monthly Weather Review* 91.3, pp. 99–164.

Sørensen, Jens Nørkær and Asger Myken (1992). "Unsteady Actuator Disc Model for Horizontal Axis Wind Turbines". In: *Journal of Wind Engineering and Industrial Aerodynamics* 39.1, pp. 139–149.

Sørensen, Jens Nørkær and Wen Zhong Shen (1999). "Computation of Wind Turbine Wakes Using Combined Navier-Stokes/Actuator-line Methodology". In: *Proceedings of European Wind Energy Conference EWEC '99, Nice*, pp. 156–159.

Stanfel, P., K. Johnson, C. J. Bay, and J. King (2021). "Proof-of-Concept of a Reinforcement Learning Framework for Wind Farm Energy Capture Maximization in Time-Varying Wind". In: *Journal of Renewable and Sustainable Energy* 13.4, p. 043305.

Stoll, Rob, Jeremy A. Gibbs, Scott T. Salesky, William Anderson, and Marc Calaf (2020). "Large-Eddy Simulation of the Atmospheric Boundary Layer". In: *Boundary-Layer Meteorology* 177.2, pp. 541–581.

Sun, Haiying, Changyu Qiu, Lin Lu, Xiaoxia Gao, Jian Chen, and Hongxing Yang (2020). "Wind Turbine Power Modelling and Optimization Using Artificial Neural Network with Wind Field Experimental Data". In: *Applied Energy* 280, p. 115880.

Sutton, Richard S. and Andrew G. Barto (2018). *Reinforcement Learning: An Introduction*. Second edition. Adaptive Computation and Machine Learning Series. Cambridge, Massachusetts: The MIT Press.

Taschner, E., AAW van Vondelen, R. Verzijlbergh, and JW van Wingerden (2023). "On the Performance of the Helix Wind Farm Control Approach in the Conventionally Neutral Atmospheric Boundary Layer". In: *Journal of Physics: Conference Series* 2505.1, p. 012006.

Troldborg, Niels (2009). "Actuator Line Modeling of Wind Turbine Wakes". In:

Veers, Paul, Katherine Dykes, Eric Lantz, Stephan Barth, Carlo L. Bottasso, Ola Carlson, Andrew Clifton, Johney Green, Peter Green, Hannele Holttinen, Daniel Laird, Ville Lehtomäki, Julie K. Lundquist, James Manwell, Melinda Marquis, Charles Meneveau, Patrick Moriarty, Xabier Munduate, Michael Muskulus, Jonathan Naughton, Lucy Pao, Joshua Paquette, Joachim Peinke, Amy Robertson, Javier Sanz Rodrigo, Anna Maria Sempreviva, J. Charles Smith, Aidan Tuohy, and Ryan Wiser (2019). "Grand Challenges in the Science of Wind Energy". In: *Science* 366.6464, eaau2027.

Verstappen, Roel (2011). "When Does Eddy Viscosity Damp Subfilter Scales Sufficiently?" In: *Journal of Scientific Computing* 49.1, p. 94.

Vignon, Colin, Jean Rabault, and Ricardo Vinuesa (2023). "Recent Advances in Applying Deep Reinforcement Learning for Flow Control: Perspectives and Future Directions". In: *Physics of Fluids* 35.

Williams, Ronald J. (1992). "Simple Statistical Gradient-Following Algorithms for Connectionist Reinforcement Learning". In: *Machine Learning* 8.3, pp. 229–256.

Xu, Hui and Pierre Sagaut (2013). "Analysis of the Absorbing Layers for the Weakly-Compressible Lattice Boltzmann Methods". In: *Journal of Computational Physics* 245, pp. 14–42.