WIND ENERGY SCIENCE

Open Access

JHTDB-wind: a web-accessible large-eddy simulation database of a wind farm with virtual sensor querying

Xiaowei Zhu[1], Shuolin Xiao[2], Ghanesh Narasimhan[3,4], Luis A. Martinez-Tossas[5], Michael Schnaubelt[6], Gerard Lemson[6], Hanxun Yao[6,7], Alexander S. Szalay[6,8,9], Dennice F. Gayme[2,6,7], and Charles Meneveau[2,6,7]

[1]Department of Mechanical and Materials Engineering, Portland State University, Portland, OR 97201, USA
[2]Ralph O'Connor Sustainable Energy Institute, Johns Hopkins University, Baltimore, MD 21218, USA
[3]Department of Mechanical Engineering, University of Minnesota, Minneapolis, MN 55455, USA
[4]St. Anthony Falls Lab, University of Minnesota, Minneapolis, MN 55455, USA
[5]National Renewable Energy Laboratory, Golden, CO 80401, USA
[6]Institute for Data Intensive Engineering and Science, Johns Hopkins University, Baltimore, MD 21218, USA
[7]Department of Mechanical Engineering, Johns Hopkins University, Baltimore, MD 21218, USA
[8]Department of Physics and Astronomy, Johns Hopkins University, Baltimore, MD 21218, USA
[9]Department of Computer Science, Johns Hopkins University, Baltimore, MD 21218, USA

Correspondence: Charles Meneveau (meneveau@jhu.edu)

Abstract. This paper introduces JHTDB-wind (https://turbulence.idies.jhu.edu/datasets/windfarms, last access: 11 November 2025), a publicly accessible database containing large-eddy simulation (LES) data from wind farms. Building on the framework of the Johns Hopkins Turbulence Database (JHTDB), which hosts direct numerical simulation (DNS) and some LES datasets of canonical turbulent flows, JHTDB-wind stores the 4D space–time history of the flow and provides users the ability to access and query the data via a web-based virtual sensor interface. The initial dataset comprises LES results from a large wind farm with 10×6 turbines, modeled using a filtered actuator line method, under conventionally neutral atmospheric conditions. These data comprise 1 h (hour) of flow field data (velocity, pressure, potential temperature deviation, subgrid-scale (SGS) eddy viscosity, and turbine forces, approximately 15 TB (terabytes) and wind turbine data – including both turbine-level operational quantities and blade-level aerodynamic quantities (approximately 1.3 TB) – stored in Zarr and Parquet formats, respectively. Data retrieval is facilitated by the giverny Python package, allowing remote users to query the database in Python or MATLAB (C and Fortran support are available for flow field data). This paper details the simulation setup and demonstrates data access through examples that analyze wind farm flow structures and turbine performance. The framework is extensible to future datasets, including the JHTDB-wind diurnal cycle simulation analyzed in Xiao et al. (2025).

1   Introduction

Eddy-resolving simulations of atmospheric boundary layer (ABL) phenomena (Porté-Agel et al., 2000; Bou-Zeid et al., 2004; Kumar et al., 2006) and of wind farms in particular (Calaf et al., 2010; Meyers and Meneveau, 2012; Gebraad et al., 2016; Stevens and Meneveau, 2017; Zhang et al., 2023) have significantly advanced our understanding of the complex, multi-scale, and multi-physics processes involved. Large-eddy simulations (LESs) offer high spatial and temporal resolution, capturing the dynamics of relatively small and fast turbulent eddies (Churchfield et al., 2012; Chatelain et al., 2013; Yang et al., 2021; Li et al., 2022). While the range of resolved scales in LES is constrained by computational resources, the number of LES grid points in typical simulations continues to increase. However, data handling and post-processing capabilities have not kept pace with the resulting rapid increase in data volumes. For instance, a single LES of turbulent flow outputting five field variables (e.g., the three velocity components, potential temperature, and pressure) on 2048^3 spatial grid points and integrated over, say, 10^4 time steps (McWilliams et al., 1994; Alexakis et al., 2024) can generate petabytes (PB) of data. As a result, most studies store only a few selected snapshots and rely heavily on pre-defined run-time diagnostics when time-resolved analysis is required. This approach reduces storage requirements but limits the ability to revisit data when new questions and concepts arise, often necessitating costly recomputation. Furthermore, certain analyses – such as backward-in-time particle tracking from an extreme dissipation event – cannot be performed without the full temporal data.

To address these challenges, modern database technologies have increasingly been applied to preserve and store data from simulation-based turbulence research (Perlman et al., 2007; Zhang et al., 2018; Chung et al., 2022; Duraisamy et al., 2019). One example is the Johns Hopkins Turbulence Database (JHTDB; https://turbulence.idies.jhu.edu, last access: 11 November 2025), an open-access platform supported by the National Science Foundation (Perlman et al., 2007; Li et al., 2008). JHTDB enables researchers to interact with easily accessible large-scale simulation data. The system currently hosts more than 1 PB of direct numerical simulation (DNS) data for canonical, turbulent flows of fundamental interest (over 2 PB if counting warm backup copies), including six space–time-resolved datasets and several others with a few snapshots available. Some LES datasets of stably stratified atmospheric turbulence are also included in JHTDB. Through web-service-based tools, users can query the database using a "virtual sensors" interface, specifying spatial and temporal locations for which the system returns properly interpolated field or derivative values (Li et al., 2008; Yu et al., 2012). A hallmark of the platform is that it allows users to access only the specific subsets of the data they require, eliminating the need to download massive datasets or manage complex file formats. This approach

has significantly broadened access to high-fidelity eddy-resolving simulation data and has contributed to democratizing high-performance computational turbulence research. To date, JHTDB data have been used in research reported in over 400 peer-reviewed journal articles.

At the same time, with the growing global demand for renewable energy, enhancing wind energy efficiency has become a key priority. As wind turbines grow larger and wind farms expand in scale, their interactions with the ABL become increasingly complex – particularly with respect to wake dynamics, energy extraction, and the redistribution of momentum within the flow. LESs of large wind turbines have emerged as a crucial complement to field measurements, enabling researchers to explore flow–turbine interactions in detail and to develop engineering models that inform turbine placement strategies and improve wind farm efficiency. For example, Calaf et al. (2010) used LES with periodic boundary conditions to study the performance of "infinite" arrays of wind turbines under neutrally stratified conditions. Abkar and Porté-Agel (2013, 2014) examined how wind farm density and free-atmosphere stability influence kinetic energy fluxes in a conventionally neutral boundary layer (CNBL) – defined as neutrally stratified surface layers capped by stably stratified free atmospheres (Zilitinkevich et al., 2002). Allaerts and Meyers (2015) explored the effect of capping inversion profile on wind farm performance. Numerous additional LES-based studies have further advanced the field (Yang et al., 2014; Aitken et al., 2014; Martínez-Tossas et al., 2015; Stevens et al., 2018; Gharaati et al., 2022, 2024; Aiyer et al., 2024), highlighting the continued value of high-resolution simulation tools for understanding and optimizing wind energy systems.

These simulations, like many previous numerical studies of large-scale wind farms, generate extensive datasets. However, access to these data often remains restricted to the original researchers who conducted the simulations. The data (typically 4D space–time fields of velocity, temperature, etc.) are ephemeral: they must be analyzed in real time during the simulation, or, at best, a limited number of snapshots are stored for post-processing, while the large majority of the data is discarded. As demonstrated in the case of the JHTDB database, providing access to the 4D space–time history of a simulation could provide substantial benefits for the broader research community. The value of open access to time-resolved numerical datasets is now being recognized beyond fluid dynamics, particularly in the field of geosciences. For example, the recently released NOW-23 dataset (Bodini et al., 2024) comprises a full year of Weather Research and Forecasting (WRF) model simulations of offshore wind conditions over several expansive (hundreds of km) US coastal regions, offering valuable data for wind farm developers. However, no equivalent open-access LES datasets currently exist at smaller scales that explicitly include wind turbine effects – datasets that would be highly valuable for researchers focused on wake interactions, tur-

bine siting, and wind farm optimization. More in general, the lack of data sharing in the wind energy sector has been recognized to hinder technical progress and leads to missed opportunities for improving the efficiency of energy markets (Kusiak, 2016).

To begin addressing the need for open access to LES wind farm data, we construct JHTDB-wind (see https://turbulence. idies.jhu.edu/datasets/windfarms, last access: 11 November 2025; Zhu et al., 2025), a publicly accessible turbulence database built on the JHTDB framework. This paper presents the dataset by detailing the simulation framework (Sect. 2) and flow configuration – specifically, a CNBL interacting with a 60-turbine wind farm using National Renewable Energy Laboratory (NREL) 5 MW reference turbines. Here, CNBL is chosen because it is a less complicated atmospheric state, observed in nature (Liu and Stevens, 2022), for example, during the transition period after sunset or on cloudy days with powerful winds (Allaerts and Meyers, 2017; Liu et al., 2024). Simulation parameters are described in Sect. 3. The construction of the database system is described in Sect. 4, followed by an overview of representative data access methods based on the JHTDB virtual sensor method, illustrated here via Python examples (Sect. 5). Conclusions are summarized in Sect. 6. Further documentation is available directly on the database website.

2 Large-eddy simulation framework

In this study, we use the open-source LES code LESGO (https://lesgo.me.jhu.edu, last access: 11 November 2025) as a numerical solver to simulate ABL flows and its interactions with wind turbines (Calaf et al., 2010; Stevens and Meneveau, 2017; Martinez, 2017; Stevens et al., 2018; Shapiro et al., 2018, 2020; Gharaati et al., 2022; Narasimhan et al., 2022, 2024a, 2025, 2024b; Gharaati et al., 2024; Ayala et al., 2024). The model represents all variables on a 3D Cartesian grid, with x, y, and z denoting the streamwise, spanwise, and vertical directions, respectively. In index notation, these are expressed as x_i, where $i = 1, 2, 3$. The corresponding velocities are denoted by u_i or also with u, v, and w for its x-, y-, and z-direction components, respectively.

2.1 Governing equations and numerical methods

The turbulent flow is simulated by solving the filtered Navier–Stokes equations in their rotational form with Boussinesq thermal forcing and Coriolis effects, along with the transport equation for the potential temperature field. The governing equations include the filtered mass conservation,

$$\frac{\partial \tilde{u}_i}{\partial x_i} = 0; \tag{1}$$

the filtered momentum conservation,

$$\frac{\partial \tilde{u}_i}{\partial t} + \tilde{u}_j \left(\frac{\partial \tilde{u}_i}{\partial x_j} - \frac{\partial \tilde{u}_j}{\partial x_i} \right) = -\frac{\partial \tilde{p}^*}{\partial x_i} + \frac{g}{\theta_0}(\tilde{\theta} - \theta_0)\delta_{i3}$$
$$- \frac{\partial \tau_{ij}^{\mathrm{SGS},d}}{\partial x_j} - f_i + f_c(\tilde{u}_2 - V_g)\delta_{i1}$$
$$- f_c(\tilde{u}_1 - U_g)\delta_{i2}; \tag{2}$$

and the filtered heat conservation,

$$\frac{\partial \tilde{\theta}}{\partial t} + \tilde{u}_j \frac{\partial \tilde{\theta}}{\partial x_j} = -\frac{\partial \Pi_j}{\partial x_j}. \tag{3}$$

Here, the tilde indicates filtering at the LES grid scale $\tilde{\Delta} = \sqrt[3]{\Delta x\ \Delta y\ \Delta z}$; ρ is the density of air; $\tau_{ij}^{\mathrm{SGS}} = \widetilde{u_i u_j} - \tilde{u}_i \tilde{u}_j$ is the subgrid-scale (SGS) stress tensor, and $\tau_{ij}^{\mathrm{SGS,d}} = \tau_{ij}^{\mathrm{SGS}} - \delta_{ij}\tau_{kk}^{\mathrm{SGS}}/3$ is the deviatoric (trace-free) part of τ_{ij}^{SGS}, where δ_{ij} is the Kronecker delta; $\tilde{p}^* = \tilde{p}/\rho + \tilde{u}_k \tilde{u}_k/2 + \tau_{kk}^{\mathrm{SGS}}/3$ is the pseudo-pressure, where $\tilde{p}$ is the resolved pressure; $g = 9.81\,\mathrm{m\,s}^{-2}$ is the gravitational acceleration; θ_0 is the reference potential temperature scale; and f_i is the distributed body force for modeling the turbine-induced aerodynamic forces on the air flow (see Sect. 2.3). In the present study, $\tau_{ij}^{\mathrm{SGS,d}}$ is parameterized using the Lilly–Smagorinsky eddy-viscosity-type model (Smagorinsky, 1963; Lilly, 1966), i.e., $\tau_{ij}^{\mathrm{SGS,d}} = -2\nu_{\mathrm{SGS}}\tilde{S}_{ij} = -2(C_s\tilde{\Delta})^2|\tilde{S}|\tilde{S}_{ij}$, where $\tilde{S}_{ij} = 0.5(\partial \tilde{u}_i/\partial x_j + \partial \tilde{u}_j/\partial x_i)$ is the resolved strain-rate tensor, $|\tilde{S}| = \sqrt{2\tilde{S}_{ij}\tilde{S}_{ij}}$ is the strain-rate magnitude, and $\nu_{\mathrm{SGS}} = (C_s\tilde{\Delta})^2|\tilde{S}|$ is the modeled SGS eddy viscosity. The coefficient C_s is dynamically determined using the Lagrangian-averaged scale-dependent dynamic model (Bou-Zeid et al., 2005), which has been successfully applied in several prior LES studies of wind turbine wake flows (Calaf et al., 2010; Stevens and Meneveau, 2017; Martinez, 2017; Stevens et al., 2018; Narasimhan et al., 2022; Gharaati et al., 2022; Narasimhan et al., 2024a; Gharaati et al., 2024). In Eq. (3), the term $\Pi_j = \widetilde{u_j\theta} - \tilde{u}_j\tilde{\theta}$ is the SGS heat flux whose eddy diffusivity (κ_{SGS}) is determined from $\kappa_{\mathrm{SGS}} = Pr_{\mathrm{SGS}}^{-1}\nu_{\mathrm{SGS}}$, where the SGS Prandtl number of $Pr_{\mathrm{SGS}} = 1$ (Narasimhan et al., 2022) is prescribed.

The atmospheric boundary layer flow is driven by a geostrophic wind whose pressure gradient is given by $-\nabla P_\infty/\rho = (f_c V_g, -f_c U_g)$. Here, $f_c = 2\Omega\sin\phi = 10^{-4}\,\mathrm{s}^{-1}$ is the Coriolis parameter corresponding to a mid-latitude position (specifically to $\phi = 43.44°$ with Earth's rotation rate $\Omega = 7.27 \times 10^{-5}\,\mathrm{rad\,s}^{-1}$). The quantities U_g, V_g are the geostrophic wind velocity components along the x and y directions, respectively, with magnitude $G = \sqrt{U_g^2 + V_g^2}$, and directed at an angle of α_G relative to the x direction such that $U_g = G\cos\alpha_G$, $V_g = G\sin\alpha_G$. At each time step, a proportional–integral (PI) controller is utilized to control the direction of the geostrophic wind such that the wind flows in the streamwise direction with zero

wind veer at the hub height (Sescu and Meneveau, 2014; Narasimhan et al., 2022).

LESGO uses a Fourier-series-based pseudo-spectral method based on collocated grids for the spatial discretizations in the horizontal (x and y) directions and a second-order central difference method based on staggered grids in the vertical (z) direction. The 3/2 rule is used to eliminate the aliasing error associated with the pseudo-spectral discretization of the nonlinear convective terms. The simulation is advanced in time using a fractional-step method. First, the velocity field is advanced in time by integrating Eq. (2) using the second-order Adams–Bashforth scheme to obtain a predicted velocity field. Then, a pressure Poisson equation is constructed based on the divergence-free constraint Eq. (1) for the new time step and is solved to obtain the pseudo-pressure field. Lastly, the predicted velocity field is projected to the divergence-free space using the gradient of the pseudo-pressure to obtain the velocity field for the new time step. The above fractional steps are repeated at every time step in LES to advance the flow field in time. More details of the numerical schemes used in the LESGO solver can be found in the original references (Albertson, 1996; Albertson and Parlange, 1999).

2.2 Boundary conditions

In the streamwise (x) direction, inflow–outflow boundary conditions are applied using the concurrent precursor simulation approach (Stevens et al., 2014). Specifically, a separate precursor domain without wind turbines is simulated to generate realistic turbulent inflow conditions, which are then imposed at the inlet of the wind farm domain. To ensure periodicity, a fringe region is introduced at the end of the wind farm domain where the outflow is gradually forced to match the inflow from the mapped region in the precursor domain. More details of the inflow–outflow conditions implemented in the current pseudo-spectral solver are provided in Stevens et al. (2014). Additionally, the simulation in the precursor domain uses a shifted periodic boundary condition where the flow field in a spanwise shifting region is shifted to prevent persistent spanwise locking of large-scale turbulent structures (Munters et al., 2016). Following the recommendation in Munters et al. (2016) a shift of $L_{y-\text{shift}} = 0.25L_z$ is used in this study, where L_z is the domain height. In the spanwise (y) direction, periodic boundary conditions are used. In the vertical (z) direction, the ground surface boundary condition is specified in both the precursor and wind turbine domains using the Monin–Obukov Similarity Theory (MOST)-based equilibrium surface flux modeling (Monin and Obukhov, 1954). The components of local surface shear stress are computed as a function of the prescribed roughness length according to

$$\tau_{i,3|\text{surf}} = -u_*^2 \frac{\widehat{\widetilde{u}}_i}{\sqrt{\widehat{\widetilde{u}}^2 + \widehat{\widetilde{v}}^2}}, \; i = 1, 2;$$

$$\text{and } u_* = \kappa \frac{\sqrt{\widehat{\widetilde{u}}^2(0.5\Delta z) + \widehat{\widetilde{v}}^2(0.5\Delta z)}}{\ln(0.5\Delta z/z_0)}. \tag{4}$$

Here, $\kappa = 0.41$ is the von Kármán constant, and z_0 is the prescribed roughness length. The friction velocity u_* is expressed in terms of the horizontal velocity $(\widehat{\widetilde{u}}, \widehat{\widetilde{v}})$ at the first grid point ($z_1 = 0.5\Delta z$), filtered at twice the grid resolution, $\widehat{\widetilde{\Delta}} = 2\widetilde{\Delta}$ (Bou-Zeid et al., 2005). Since we simulate conventionally neutral conditions, the surface heat flux is set to zero; thus no stability correction terms (as used in Xiao et al., 2025) are included. At the top of the domain, a stress-free boundary condition is imposed. A sponge or Rayleigh-damping layer (Durran and Klemp, 1983) is included approaching the top boundary, ranging from $0.75L_z$ to L_z, with a sponge inverse relaxation timescale (frequency) parameter of $3.9 \times 10^{-3}\,\text{s}^{-1}$. In this layer, a damping body force with a cosine profile is applied to suppress the reflection of gravity waves.

Henceforth, the $\widetilde{(\cdot)}$ notation for LES-filtered field variables (e.g., velocity $\widetilde{u}_i$, temperature $\widetilde{\theta}$) will be omitted for brevity. All subsequent variables should be interpreted as implicitly filtered quantities obtained from the LES solution, governed by the equations presented in this section.

2.3 Wind turbine representation

The aerodynamic forces exerted by wind turbines on the airflow are modeled through the distributed body force term f_i in the momentum transport equations (Eq. 2). During the initial spin-up phase (i.e., Phase 1), we employ an actuator disk model (ADM) on a coarse grid for computational efficiency, with the thrust force magnitude calculated as $f = \frac{\pi}{8}\rho C_T' \langle u_T \rangle_d^2 D^2$ (Calaf et al., 2010; Howland et al., 2016). Here, ρ is the air density, $\langle u_T \rangle_d$ is the local wind velocity averaged over the rotor disk, D is the diameter of the wind turbine, and C_T' is the local thrust coefficient (set to a common value $C_T' = 1.33$). We recall that C_T' is based on the disk-averaged velocity $\langle u_T \rangle_d$, which, unlike the far-upstream velocity U_∞, is immediately available in LES (Calaf et al., 2010).

After the spin-up simulation converges to quasi-steady behavior, the grid is refined to its final resolution, and the actuator line model (ALM) is adopted (Sørensen and Shen, 2002; Troldborg, 2009; Jha et al., 2014; Martínez-Tossas et al., 2015). In ALM, each turbine blade is represented by a collection of actuator points along a line, where forces are applied according to the velocity field and the angle of attack. The forces per unit width at every actuator point are computed as

$$\mathbf{f}_{\text{alm}} = 0.5\rho c |\mathbf{V}_{\text{rel}}|^2 (C_L \boldsymbol{e}_L + C_D \boldsymbol{e}_D), \tag{5}$$

where c is the airfoil chord length, $|\mathbf{V}_{\text{rel}}|$ is the magnitude of the relative velocity of the upwind flow to the turbine blade, C_L and C_D are lift and drag coefficients obtained from tabulated airfoil data, and $\boldsymbol{e}_L$ and $\boldsymbol{e}_D$ are unit vectors along the direction of the lift and drag forces at each actuator point, respectively. These forces are then smeared using a Gaussian kernel to project them into the computational LES grid:

$$\eta_\epsilon = \frac{1}{\epsilon^3 \pi^{3/2}} e^{-r^2/\epsilon^2}, \tag{6}$$

where r is the distance from the grid point to the actuator point and ϵ denotes the width of the kernel. The kernel width is chosen to be at least $\epsilon = 2(\Delta_x \Delta_y \Delta_z)^{1/3}$, as recommended to avoid numerical instabilities (Troldborg, 2009; Martínez-Tossas et al., 2015).

The accuracy of the ALM can be sensitive to grid resolution and the choice of ϵ. The optimal ϵ_{opt} needed to resolve the induced velocities is typically much smaller than the ϵ used to avoid numerical instabilities (Martínez-Tossas et al., 2017). To address this challenge, we use the generalized filtered lifting line theory correction to accurately represent the blade aerodynamics (Martínez-Tossas and Meneveau, 2019; Martínez-Tossas et al., 2024), including the shedding of unresolved vorticity leading to missing induced velocities at the blade. The correction accounts for subgrid-scale induced velocity that would be obtained by using an optimal ϵ_{opt} by estimating its contribution and adding it to the resolved velocity in the LES. With the correction, the ALM provides consistent blade loading predictions across varying grid resolutions.

The NREL-5MW baseline wind turbine (Jonkman et al., 2009) is adopted as our reference model. It is a widely used benchmark model developed by NREL to standardize research on wind technologies. The turbine has a diameter of $D = 126$ m, three blades, and a hub height at elevation $z_h = 90$ m. It reaches a rated electrical power output of 5 MW at a rated wind speed of approximately $11.4 \, \text{m s}^{-1}$. Its rotor blades utilize the Delft University (DU) and National Advisory Committee for Aeronautics (NACA) series airfoil profiles optimized for aerodynamic efficiency, structural integrity, and minimal fatigue loads, making the NREL-5MW turbine an essential tool for evaluating wind turbine performance, control strategies, structural design, and off-shore platform dynamics.

The dataset employs fixed but row-dependent rotor angular velocities determined through an initialization procedure. Initialization begins with all turbines operating at tip-speed ratio $\text{TSR} = 7.5$ (near-optimal for NREL-5MW turbines). In this initialization simulation (i.e., first part of Phase 2), the angular velocity Ω for each turbine is then computed dynamically using

$$\Omega = \text{TSR} \times \frac{1.087 \, U_d}{(1-a) R}, \tag{7}$$

where U_d is the disk-averaged velocity; the numerator incorporates an empirical 8.7 % correction factor for LES filter-

scale effects ($\epsilon = 16$ m), validated through single-turbine laminar inflow tests; and the induction factor a derives from rotor geometry (blade number $N_b = 3$, radius $R = 63$ m, and chord $c = 3$–4 m) and local inflow angle ϕ via

$$a = \frac{1}{(4\sin^2\phi)/(\sigma_r C_n) + 1}, \tag{8}$$

with a rotor solidity of $\sigma_r = N_b c/(\pi R)$ and a force coefficient of $C_n = C_L \cos\phi + C_D \sin\phi$. After approximately 40 min (minutes) of initialization simulation, the angular velocity Ω for each turbine is averaged within its respective row, which serves as the fixed operational values for the subsequent database simulations.

We also note that LESGO's ALM implementation includes detailed turbine operation control methods, such as pitching the blades (feathering) during region 3 operations. In the current simulation we chose to operate all turbines exclusively at optimal tip-speed ratio, "region 2" (also without including regions 1.5 and 2.5). This choice was made in order to avoid the need to store additional data relating to blade pitch (curtailment) and other complex turbine control actions. Since this practice deviates slightly from the reference NREL-5MW nameplate data, we refer to the turbine in our simulations as the NREL-5MW+ turbine. Specifically, the front turbines are allowed to rotate slightly faster than the maximum rotation rate of the original NREL-5MW reference turbine.

3 Simulation parameters

We simulate turbulent flow through a 10×6 array of NREL-5MW+ turbines (with diameter $D = 126$ m) in a $28.224 \times 3.78 \times 2 \, \text{km}^3$ domain, equally split between precursor and wind farm subdomains (each $112D = 14.112$ km long). Figure 1 displays the domain dimensions. The precursor domain includes the region denoted as P of length $5L_x/8$, the mapping region P_M of length $L_x/8$, and the spanwise shifting region P_S of length $L_x/8$. The wind farm domain features $14D$ of upstream buffer zone, a $63D$ turbine region, a $21D$ downstream wake recovery region (these three regions combined are denoted as W), and a $14D$ outflow fringe region (W_F). The turbines are spaced $7D$ (streamwise) and $5D$ (spanwise), with lateral boundaries $2.5D$ from the outermost turbines. Note that P_M, P_S and W_F regions have a length of $L_x/8$ and that P_M extends from $5L_x/8$ to $3L_x/4$. Vertically, a 0.5 km Rayleigh damping sponge layer (denoted as R) is located between 1.5 and 2 km (see Fig. 1). We adopt $\theta_0 = 263.5$ K as the reference potential temperature, consistent with the value chosen in studies by Gadde and Stevens (2021) and our prior simulations of stable boundary layer (SBL) and CNBL flows reported in Narasimhan et al. (2024a). This reference temperature was inspired by observations from the Beaufort Sea Arctic Stratus Experiment (BASE) and simulations by Kosović and Curry (2000). While the value of θ_0 is relatively low,
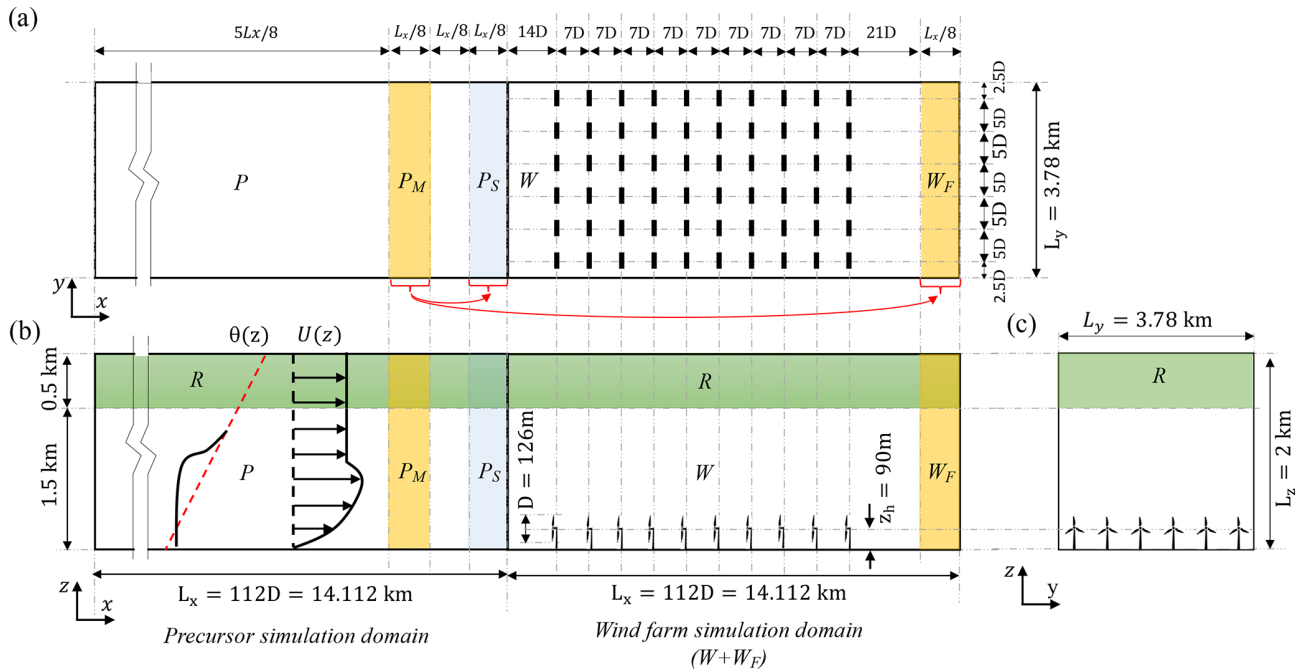
Figure 1. Schematic representation of the computational simulation domain (not to scale), showing (**a**) top view (x–y plane), (**b**) side view (x–z plane), and (**c**) front view (y–z plane). The precursor computational domain consists of the regions denoted as "P", the precursor mapping region "P_M", and the precursor spanwise shifting region "P_S". The wind farm computational domain includes the wind farm region "W" and the fringe region "W_F" near the outlet. Both the precursor and wind farm computational domains include a Rayleigh damping region at the top (denoted as "R"). The turbine diameter $D = 126$ m and hub height $z_h = 90$ m are also marked.

it serves primarily as a relative additive reference that does not significantly affect the simulated flow dynamics or the physical interpretation of the results. For example, if we used 273 K, it would change the implied thermal expansion coefficient in our Boussinesq approximation only by about 3 %.

The turbulent flow is driven by a constant geostrophic wind speed $G = 15\,\mathrm{m\,s^{-1}}$ at $\alpha_g \approx -22.5°$ to the x direction, with the angle controlled by a PI controller ($K_P = 10$, $K_I = 0.5$) to align hub-height mean wind velocity with the x axis in the conventionally neutral boundary layer (Sescu and Meneveau, 2014; Narasimhan et al., 2022). The surface has roughness length of $z_0 = 0.1$ m and a reference potential temperature of $\theta_0 = 263.5$ K. Initial conditions are set at $U_g = 15\,\mathrm{m\,s^{-1}}$ (streamwise) and $V_g = 0\,\mathrm{m\,s^{-1}}$ (spanwise), perturbed by random noise, while potential temperature decreases from 265 K at the surface with a $1\,\mathrm{K\,km^{-1}}$ lapse rate, including random perturbations below 1 km.

The numerical simulation is conducted in three consecutive phases to ensure proper flow development and statistical convergence.

- Phase 1: coarse-resolution ADM spin-up. A 10 h simulation using the ADM is performed to establish a quasi-stationary atmospheric boundary layer and wind farm wake field. This phase leverages the computational efficiency of ADM, which approximates turbine forces without resolving actuator line-level aerodynamics.

- Phase 2: fine-resolution ALM convergence. A 1 h simulation using the actuator line model at finer spatial resolution transitions the flow from ADM-averaged to ALM-resolved turbine representation. Besides the turbine model update, two additional changes are introduced in this phase: (i) the time-stepping scheme is switched from a constant Courant–Friedrichs–Lewy (CFL) number of 0.0625 to a fixed time step of $\Delta t = 0.025$ s. This adjustment has negligible impact on the results because, under these simulation conditions, CFL = 0.0625 corresponds to $\Delta t \approx 0.03$ s. The slightly more restrictive $\Delta t = 0.025$ s maintains numerical stability while preserving solution accuracy. (ii) The rotor control changes from a fixed tip-speed ratio (TSR = 7.5) to fixed rotor angular velocities that vary across turbine rows, as tabulated in Table 1. This adjustment has a negligible impact on the results because the prescribed angular velocities closely match the values achieved under TSR = 7.5 conditions (see the calculation method in Sect. 2.3), ensuring nearly identical rotor dynamics.

- Phase 3: fine-resolution simulation for database construction. A final 1 h simulation is carried out to collect high-fidelity flow and turbine data. Flow field variables are recorded every 20 LES time steps (i.e., every 0.5 s) on a filtered and subsampled spatial grid (every other grid point in the x–y plane), while wind turbine data

– both integral and blade-resolved – are stored at every LES time step (0.025 s). Note that we purposefully operate the NREL-5MW+ turbine in "region 2" during the simulation time, in order to avoid having to choose and document additional controller actions. As a result, during some times, some of the turbines operate "above rated conditions" but maintain self-consistent aerodynamic behavior of the blades and air flow.

The three phases of the simulation are illustrated through the time history of the boundary layer height $z_i = h_{\text{ABL}}$ and the geostrophic wind angle shown in Fig. 2.

4 JHTDB-wind database construction

The LES data from the final 1 h sampling period are systematically ingested into the database and organized into two primary data types: (i) flow field data, consisting of 4D space–time fields captured across both simulation domains (precursor and wind farm domains), providing complete spatiotemporal information about the atmospheric flow, and (ii) turbine data, which are further subdivided into two subtypes. The first subtype is turbine-level operational data, comprising time histories of turbine power and thrust. The second subtype is blade-level data, which include time histories of aerodynamic quantities sampled at each discrete actuator point along each blade.

4.1 Flow field data

4.1.1 Domain of the dataset

As described in Sect. 3, the LES is conducted in the domain of dimensions $(2 \times 14.112) \times 3.78 \times 2 \, \text{km}^3$ (see Table 2). When compiling the database, we exclude numerically imposed auxiliary regions: specifically, the final $L_x/4$ of the precursor domain (which includes the precursor spanwise shifting region P_S) and the final $L_x/8$ of the wind farm domain (i.e., the wind farm fringe region W_F), as visualized in Fig. 1. These regions serve purely numerical functions (periodicity enforcement and inflow recycling, respectively) without contributing to physical flow dynamics of interest. The resulting database domain has the extents of $(10.584 + 12.348) \times 3.78 \times 2 \, \text{km}^3$, as shown in Fig. 3. The top 0.5 km sponge region is kept in the database for simplicity of data management and possible interest.

4.1.2 Spatial resolution of the dataset

To minimize storage, we applied spectral filtering on x–y planes for flow field data by truncating Fourier modes above $\kappa_{\text{max}}/2$, where $\kappa_{\text{max}} = \pi/\Delta_{\text{LES}}$ is the LES cutoff wavenumber. The filtered fields were then subsampled at every alternate grid point in the x and y directions, maintaining the original vertical (z) resolution. This approach reduces the dataset

size by 75% while maintaining fidelity in capturing the dynamically significant larger-scale flow structures and turbine wake interactions. Thus, the flow field data have a grid size of $(576 + 672) \times 192 \times 400$.

4.1.3 Temporal resolution of the dataset

Field data are stored at intervals of 0.5 s (i.e., every 20 LES steps of 0.025 s), ensuring that fluid parcels advected at the maximum geostrophic speed ($15 \, \text{m s}^{-1}$) travel less than the horizontal grid spacing ($\Delta x \approx 9.19 \, \text{m}$) between snapshots. Although rotor blade tips move across several vertical grid spacings during this interval, the corresponding rotor force field is smooth (Gaussian filtered at scale $\epsilon = 16 \, \text{m} > 2\sqrt[3]{\Delta x \, \Delta y \, \Delta z}$), ensuring that the storage frequency of 0.5 s remains appropriate. Over the 1 h simulation period (i.e., 3600 s), the simulation advances through $3600/0.025 = 144\,000$ LES time steps, with flow fields stored at $144\,000/20 = 7200$ consecutive snapshots.

4.1.4 Final structure of the dataset

Consequently, the final stored data dimensions are $n_x \times n_y \times n_z \times n_{\text{t}} = 1248 \times 192 \times 400 \times 7200$. At each stored time step, six spatial fields are recorded: the three velocity components $u(x, y, z, t)$, $v(x, y, z, t)$, and $w(x, y, z, t)$; the (kinematic) pressure field $p(x, y, z, t)/\rho = p^*(x, y, z, t) - u_k u_k/2$ (the SGS stress trace is not available and is anyhow negligible); the potential temperature field relative to the reference temperature $\theta'(x, y, z, t) = \theta(x, y, z, t) - \theta_0$; and the subgrid-scale eddy viscosity $\nu_{\text{SGS}}(x, y, z, t)$. In addition, the three components of the turbine force field, $f_x(x, y, z, t)$, $f_y(x, y, z, t)$, and $f_z(x, y, z, t)$, are also stored. Unlike the other flow field variables, these force components are stored only from the ground up to 200 m in the vertical direction. However, they are retained at the original spatial resolution (i.e., not filtered in the x–y planes). The detailed information of these stored field variables can be found in Table 3. It also needs to be mentioned that the concurrent precursor method ensures smooth transitions in velocity, potential temperature, and eddy viscosity fields between precursor and wind farm subdomains, by construction. However, due to the non-local nature of the pressure solution (solved separately in each domain via Poisson equations) and the velocity-only coupling between domains, the stored pressure field exhibits a minor discontinuity at the interface. This artifact does not affect the resolved turbulence dynamics or turbine wake interactions but needs to be taken into account if computing pressure gradients across the boundary separating the precursor and wind farm domains.

These 4D field variables are stored using Zarr format (Miles et al., 2023). In Zarr-based storage, data are organized into chunks, the smallest units retrieved during a query. To ensure efficient data access, chunk sizes must be large enough to support common operations, such as differenti-

Table 1. Rotor speed for each row of turbines.

Row no.	1	2	3	4	5	6	7	8	9	10
Ω (rad s^{-1})	1.33	1.02	1.04	1.07	1.09	1.09	1.09	1.09	1.09	1.10

Table 2. Three consecutive phases and computational domain parameters

Phase	Grid level	Turbine model	Domain size $(2 \times L_x) \times L_y \times L_z$ (km $\times$ km $\times$ km)	Number of grid points $(2 \times N_x) \times N_y \times N_z$	Spatial resolution $\Delta x \times \Delta y \times \Delta z$ (m $\times$ m $\times$ m)	Time grid CFL or Δt (– or s)
1	Coarse	ADM	$(2 \times 14.112) \times 3.78 \times 2$	$(2 \times 512) \times 192 \times 400$	$27.56 \times 19.69 \times 5$	CFL $= 0.0625$
2	Fine	ALM TSR $= 7.5$	$(2 \times 14.112) \times 3.78 \times 2$	$(2 \times 1536) \times 384 \times 400$	$9.19 \times 9.84 \times 5$	CFL $= 0.0625$
		ALM $\Omega = $ const				$\Delta t = 0.025$ s
3	Fine	ALM $\Omega = $ const	Simulation with			
			$(2 \times 14.112) \times 3.78 \times 2$	$(2 \times 1536) \times 384 \times 400$	$9.19 \times 9.84 \times 5$	$\Delta t = 0.025$ s
			Sampling over/with			
			$(10.584 + 12.348) \times 3.78 \times 2$	$(576 + 672) \times 192 \times 400$	$18.38 \times 19.68 \times 5$	$\Delta t = 0.5$ s

ations and interpolations, which typically require access to a 3D neighborhood around the query point, while remaining small enough to avoid excessive memory usage. Based on extensive testing and prior experience with other JHTDB datasets, a chunk size of 64^3 grid points provides optimal retrieval speeds and performance for typical data access modalities. We chose a similar chunk size but shaped according to $52 \times 64 \times 80$ so that an integer multiple of the chunk size in each direction fits into the stored domain size. The total amount of data stored is about 15 TB. These flow field data can be queried using `getData(...)` calls from analysis programs such as Python, MATLAB, Fortran, or C, in the same manner as with other turbulence datasets available through JHTDB.

4.2 Wind turbine data

4.2.1 Turbine-level data

The turbine-level data are integral quantities characteristic of each turbine operation, which are derived from the actuator line modeling. This dataset includes high-fidelity time histories of power output, thrust force, and rotor angular velocity, sampled at $\Delta t = 0.025$ s for all 60 turbines, as summarized in Table 4. In the present dataset, the angular velocity is held constant in time, but for other datasets (e.g., Xiao et al., 2025), this is not generally the case. For each variable, the dataset consists of 144 000 rows and 2 columns, where the first column represents time and the second column contains the corresponding values of the recorded variable. The turbine data are stored in files using Parquet format, which fa-

cilitates efficient access and querying from various programming languages. Turbine-level data can be accessed using the `getTurbineData(...)` function call from analysis environments such as Python or MATLAB.

Table 4 summarizes the turbine-level data variables. Note that, unlike the field data which are stored in kinematic (density-independent) units, the force and power data require a specified air density. The value used in the simulations to compute these forces is $\rho_{\text{air}} = 1.23$ kg m^{-3}.

4.2.2 Blade-level data

In addition to the integral quantities characteristic of each turbine's operation, more detailed information is captured along each turbine blade to enable blade-resolved aerodynamic analysis. This fine-grained dataset allows users to investigate the local aerodynamic behavior of blades under unsteady flow conditions, which is critical for understanding load distributions, fatigue effects, and control optimization strategies. The turbine blade-level dataset includes high-fidelity time histories sampled at 0.025 s for all 180 blades in the wind farm (i.e., 60 turbines $\times$ 3 blades each), with aerodynamic and geometric quantities sampled at 100 discrete actuator line points along the blade span. As summarized in Table 5, a total of 19 variables are sampled and stored, with each variable written to a separate file. For each variable, the dataset has dimensions of 144 000 $\times$ 3 rows and 103 columns. Each time step includes three rows corresponding to the three blades of a turbine, resulting in a total of 144 000 $\times$ 3 rows. Vertically, the first column represents time in seconds, the
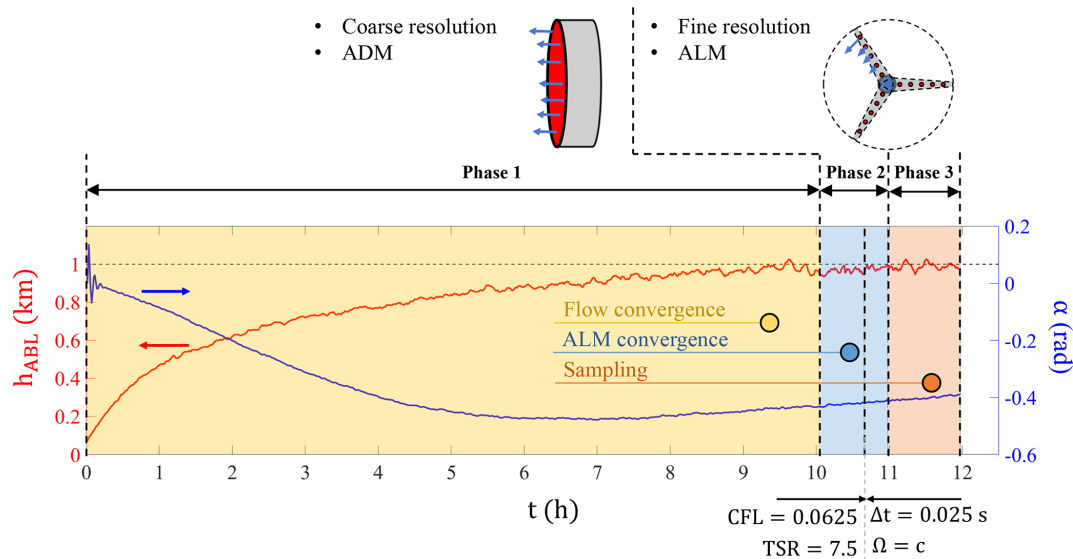
Figure 2. Time history of boundary layer height $z_i = h_{ABL}$ and geostrophic wind angle α, indicating the three simulation phases (Phase 1: coarse-resolution ADM spin-up; Phase 2: fine-resolution ALM convergence; and Phase 3: fine-resolution simulation for database construction).
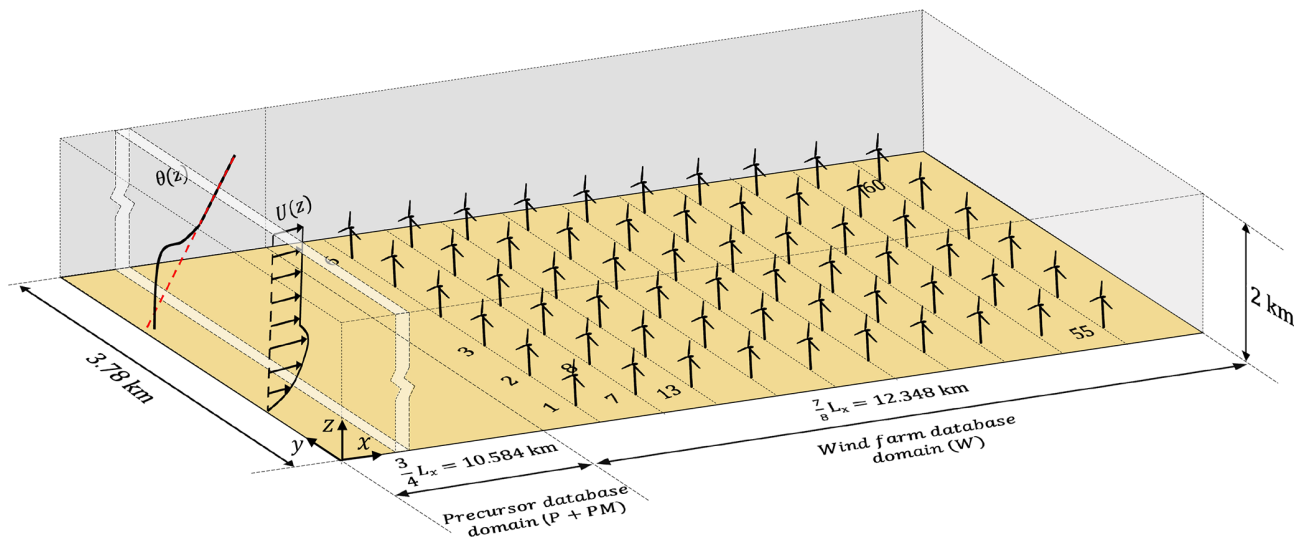

Figure 3. Schematic representation of the database domain (not to scale). This is the physical domain available in the database, merging the precursor domain ($P + P_M$) up to the end of the mapping region at $3/4 L_x$ with the wind farm domain (W) and excluding the fringe region (W_F). A total of 60 turbines are shown, with only a subset labeled for clarity. The domain dimensions are $(10.584 + 12.348) \times 3.78 \times 2\,\mathrm{km}^3$.

second column specifies the turbine number, and the third column denotes the blade number (blades can be identified by the time histories of the individual ALM point positions). The remaining 100 columns contain the values of the selected variables at each of the 100 actuator points from the blade root to tip. Similarly to turbine-level data, blade-level data are stored as Parquet files, allowing efficient access across multiple programming environments. Blade-level data can be accessed using the `getBladeData(...)` function call from analysis environments such as Python or MATLAB.

5   Web-accessible virtual sensor data access methods and examples

5.1   Flow field data

A defining feature of the JHTDB database system (Li et al., 2008) is its low entry barrier for data usage, enabling users to efficiently explore large-scale simulation datasets through web services and virtual sensor methodology. The JHTDB-wind system adopts the same approach, allowing access to wind farm data using these established tools. Users can de-

Table 3. Summary of flow field data variables.

No.	Name of variable	Name in dataset	Symbol	Unit	Data size $n_x \times n_y \times n_z \times n_\mathrm{t}$	Data resolution $\Delta x \times \Delta y \times \Delta z \times \Delta t$ $(\mathrm{m} \times \mathrm{m} \times \mathrm{m} \times \mathrm{s})$
1	Streamwise velocity		u			
2	Spanwise velocity	velocity	v	$\mathrm{m\,s^{-1}}$		
3	Vertical velocity		w		$1248 \times 192 \times 400 \times 7200$	$18.38 \times 19.68 \times 5 \times 0.5$
4	Potential temperature deviation	temperature	θ'	K		
5	Pressure (kinematic)	pressure	p	$\mathrm{m^2\,s^{-2}}$		
6	SGS eddy viscosity	eddyviscosity	ν_{SGS}	$\mathrm{m^2\,s^{-1}}$		
7	Turbine streamwise force (kinematic)		f_x			
8	Turbine spanwise force (kinematic)	force	f_y	$\mathrm{m\,s^{-2}}$	$871 \times 384 \times 40 \times 7200$	$9.19 \times 9.84 \times 5 \times 0.5$
9	Turbine vertical force (kinematic)		f_z			

Figure 4. Contour plots of instantaneous flow field variables in part of the precursor domain (here between $x = 0\,\mathrm{m}$ and $x = 10\,381.875\,\mathrm{m}$), at time $t = 1800.75\,\mathrm{s}$. **(a)** The streamwise velocity u, **(b)** the vertical velocity w, **(c)** the pressure p, and **(d)** the potential temperature deviation θ'.

velop analysis scripts or notebooks in familiar programming languages such as Python and MATLAB (as well Fortran and C) to run them remotely on their own machines or on SciServer, a cloud service dedicated to running code close to the data. Within these analysis environments, users specify space–time arrays by defining spatial locations (e.g., along a line, across a surface, within a subvolume, or scattered arbitrarily) and corresponding time instances; i.e., users specify the positions of virtual sensor arrays. These space–time arrays are then passed to the pre-defined function, `getData(...)`, which returns interpolated values of the selected variables at defined coordinates. This framework enables targeted, on-demand data access without the need to download large volumes of raw simulation output.

Figures 4 and 5 display contour plots of flow field variables at the turbine hub height ($z = z_\mathrm{h} = 90\,\mathrm{m}$) for the precursor and wind farm domains, respectively. Figure 6 presents Python code snippets that demonstrate how to query the JHTDB-wind database to extract snapshots of velocity, pressure, and potential temperature fields at a specific time, ap-

proximately in the middle of the stored 1 h dataset, namely at $t = 1800.75\,\mathrm{s}$. As a first step, an array "points" is populated with spatial coordinates that define a 2D plane: in this case, an equally spaced grid of 950×200 points in the x and y directions at a constant height $z = z_\mathrm{h} = 90\,\mathrm{m}$. These query points typically do not coincide with the actual simulation grid points, and users are not required to know the grid layout to access the data. The JHTDB-wind interface provides interpolated field values based on a user-specified interpolation method. Supported options include no interpolation (it returns the value at the nearest grid point); Lagrange polynomials of order 4, 6, or 8; and several spline interpolation methods (Li et al., 2008; Graham et al., 2016). In this example, we use eighth-order Lagrange polynomial interpolation in space. Similarly, if the requested time does not coincide with a stored time step, temporal interpolation is applied using the third-order Piecewise Cubic Hermite Interpolating Polynomial (PCHIP) method (Li et al., 2008). This user-friendly data access model eliminates the need for downloading and parsing simulation files. Instead, the Python application pro-

Table 4. Summary of turbine-level data variables. Each dataset is a 2D matrix of size $n_t \times 2$, where n_t is the number of time steps. Columns 1 and 2 represent time and measured values, respectively.

No.	Name of variable	Name in dataset	Symbol	Unit	Data size $n_t \times 2$	Data resolution Δt (s)
1	Power	power	P	W		
2	Thrust force	thrust	F_t	N	$144\,000 \times 2$	0.025
3	Rotor angular velocity	RotSpeed	Ω	$\mathrm{rad\,s^{-1}}$		

Table 5. Summary of blade-level data variables. Each dataset is a 2D matrix of size $(n_t \times 3) \times 103$. Here, $(n_t \times 3)$ represents the total number of blade-wise samples, formed by concatenating the time series data from each of the three blades of a turbine. Columns 1–3 represent time and turbine number, and columns 4–103 store aerodynamic measurements at $n_\ell = 100$ discrete locations along each blade.

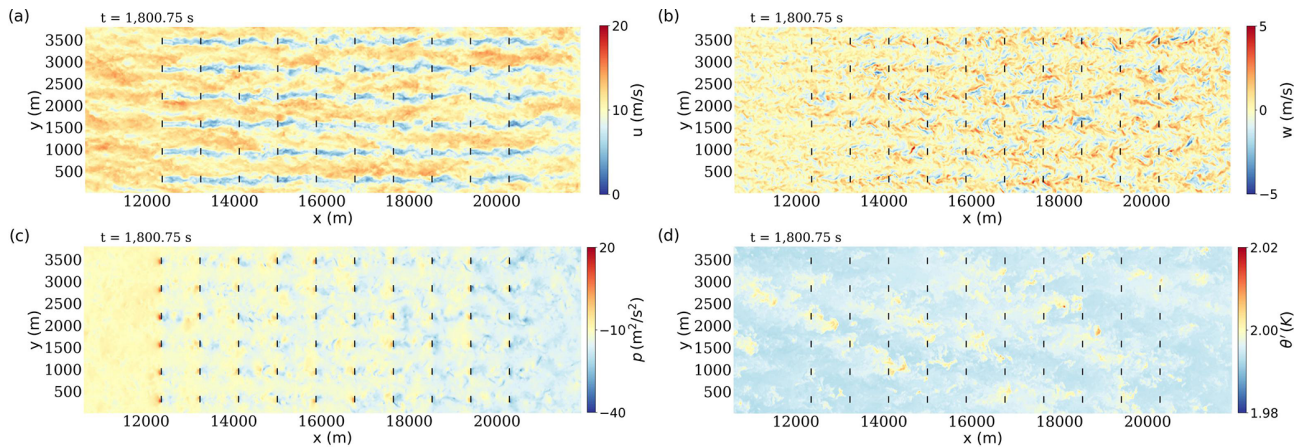| No.	Name of variable	Name in dataset	Symbol	Unit	Data size $(n_t \times 3) \times (n_\ell + 3)$	Data resolution $\Delta t \times \Delta \ell$ (s × m)
1	x position of ALM point	xPos	P_x			
2	y position of ALM point	yPos	P_y	m		
3	z position of ALM point	zPos	P_z			
4	Perturbation velocity at LES resolution, component 1	uy_LES1	$u'_{y,\mathrm{LES1}}$			
5	Perturbation velocity at LES resolution, component 2	uy_LES2	$u'_{y,\mathrm{LES2}}$	$\mathrm{m\,s^{-1}}$		
6	Perturbation velocity at optimal resolution (0.25c), component 1	uy_opt1	$u'_{y,\mathrm{opt}}$		$(144\,000 \times 3) \times (100 + 3)$	0.025×0.615
7	Perturbation velocity at optimal resolution (0.25c), component 2	uy_opt2	$u'_{y,\mathrm{opt}}$			
8	Perturbation velocity correction $u'_{y,\mathrm{opt}} - u'_{y,\mathrm{LES}}$, component 1	du1	$\Delta u'_{y,1}$			
9	Perturbation velocity correction $u'_{y,\mathrm{opt}} - u'_{y,\mathrm{LES}}$, component 2	du2	$\Delta u'_{y,2}$			
10	Angle of attack	alpha	α	rad		
11	Lift coefficient	Cl	C_L			
12	Drag coefficient	Cd	C_D	–		
13	Lift force per unit length	lift	F_L/ℓ			
14	Drag force per unit length	drag	F_D/ℓ	$\mathrm{N\,m^{-1}}$		
15	Local relative velocity magnitude	Vmag	V_mag			
16	Axial component of the local relative velocity in blade-oriented coordinates	Vaxial	V_axi	$\mathrm{m\,s^{-1}}$		
17	Tangential component of the local relative velocity in blade-oriented coordinates	Vtangential	V_tan			
18	Axial component of the local force	axialForce	F_axi	N		
19	Tangential component of the local force	tangentialForce	F_tan			

Figure 5. Contour plots of instantaneous flow field variables in part of the wind farm domain (here between $x = 10\,584$ m and $x = 21\,921.375$ m), at time $t = 1800.75$ s. **(a)** The streamwise velocity u, **(b)** the vertical velocity w, **(c)** the pressure p, and **(d)** the potential temperature deviation θ'. The short black lines represent the location of wind turbines.

```
"""
initialize getData parameters (except time and points)
"""
variable1, variable2, variable3, temporal_method, spatial_method, spatial_operator = 'velocity', 'pressure', 'temperature', 'pchip', 'lag8', 'field'

"""
initialize getData times and points
"""
time, nx, ny, n_points = 1800.75, 950, 200, 950 * 200

x_start, x_end, y_start, y_end = 10584, 21921.375, 0, 3789
x_points, y_points, z_points = np.linspace(x_start, x_end, nx, dtype=np.float64), np.linspace(y_start, y_end, ny, dtype=np.float64), 90
points = np.array([[x, y, z_points] for x in x_points for y in y_points], dtype=np.float64)

"""
use the tools and processing gizmos.
"""
# process interpolation/differentiation of points.
result1 = getData(dataset, variable1, time, temporal_method, spatial_method, spatial_operator, points)
result2 = getData(dataset, variable2, time, temporal_method, spatial_method, spatial_operator, points)
result3 = getData(dataset, variable3, time, temporal_method, spatial_method, spatial_operator, points)
```

Figure 6. Python code snippet used to obtain the data to generate Fig. 5.

gramming interface (API) returns arrays with the queried field variables, which can then be visualized directly within a Jupyter notebook (or MATLAB code). This approach was used to generate Figs. 4 and 5. It is important to note that the full 1 h dataset (comprising 14 400 time steps) is available for analysis, allowing users to query any time between $t = 0$ and $t = 3600$ s. For example, Fig. 7 shows a hub-height snapshot over the entire domain at time $t = 2505$ s.

Similar queries can be made for the values, spatial gradients, and Hessians (second-order derivatives) of all variables listed in Table 3. For example, Fig. 8a and b show the turbine streamwise force field f_x and the x-direction gradient of the pressure field ($\partial p/\partial x$), respectively, on a y–z plane intersecting Row 1 (turbines #1–#6) at $x = 12\,348$ m (1764 m downstream of the wind farm domain), at time $t = 1000.013$ s. Figure 8c and d present similar results on a plane intersecting Row 9 (turbines #49–#54) at $x = 19\,404$ m (8820 m downstream of the wind farm domain) at another time $t = 2000.67$ s. These plots were generated using the Python code shown in Fig. 9. In these examples, the queried times are in-

tentionally chosen not to coincide with the stored simulation time steps, demonstrating the temporal interpolation capabilities of JHTDB-wind.

Next, we provide examples of computed mean vertical profiles of fundamental flow quantities within the precursor domain, which features standard conventionally neutral atmospheric conditions. Figure 10 shows vertical profiles of horizontal- and time-averaged mean velocities, subgrid-scale eddy viscosity, and deviations in potential temperature, all obtained by averaging in the horizontal directions and over time. The data used to produce these profiles are retrieved using the virtual sensor framework, and an example code snippet demonstrating this process is shown in Fig. 11.

5.2   Wind turbine data

Wind turbine data, including both the turbine-level and blade-level data, are considerably smaller than the 4D flow field data, and one possibility would have been to allow users to download these data directly as files. However, such an

Figure 7. Contour plot of instantaneous streamwise velocity u in the entire database domain, ranging from $x = 0$ to $x = 22\,913.625$ m, at time $t = 2505$ s. It is noted that, although the total length of the database domain is $10\,584 + 12\,348 = 22\,932$ m, the data resolution in the x direction is 18.375 m and the grid points are located at cell centers. Consequently, the last data point is located at $22\,932 - 18.375 = 22\,913.625$ m. The short black lines represent the location of wind turbines.
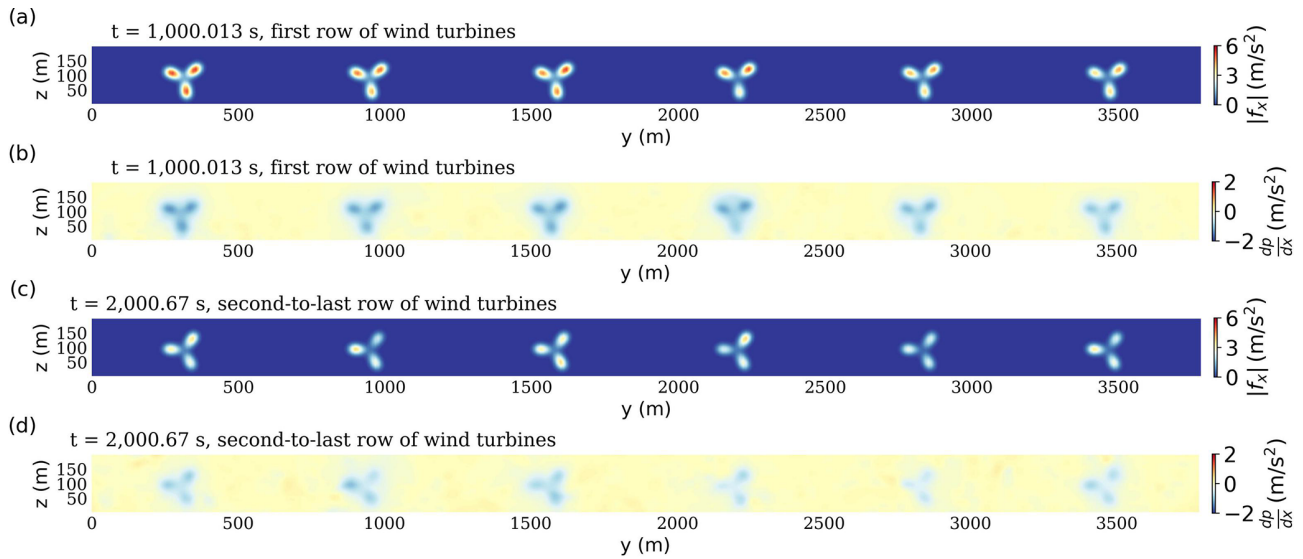

Figure 8. Instantaneous contours of turbine streamwise (i.e., x-component) force (as projected onto the LES grid using Gaussian smoothing as part of the ALM method) in the y–z planes **(a)** in the first row (i.e., Row 1, $x = 12\,348$ m) and between the relevant vertical range $z \in [2.5, 200]$ m and **(c)** in the second-to-last row (i.e., Row 9, $x = 19\,404$ m). Panels **(b)** and **(d)** show the x-direction pressure gradient distributions on the same planes, coincident with the turbines.

```
"""
initialize getData parameters (except time and points)
"""
variable1, variable2, temporal_method, spatial_method1, spatial_method2, spatial_operator1, spatial_operator2 = 'force', 'pressure', 'none', 'none', 'fd4lag4', 'field', 'gradient'

"""
initialize getData times and points
"""
time1, time2, ny, nz, n_points = 1000.013, 2000.67, 384, 40, 384 * 40

y_start, y_end, z_start, z_end = 0, 3780, 2.5, 197.5

"""
first row of wind turbines locates 12348 m downstream of inlet.
second-to-last row of wind turbines locates 19404 m downstream of inlet.
"""
x_points1, y_points1, z_points1 = 12348, np.linspace(y_start, y_end, ny, dtype=np.float64), np.linspace(z_start, z_end, nz, dtype=np.float64)
points1 = np.array([[x_points1, y, z] for y in y_points1 for z in z_points1], dtype=np.float64)
x_points2, y_points2, z_points2 = 19404, np.linspace(y_start, y_end, ny, dtype=np.float64), np.linspace(z_start, z_end, nz, dtype=np.float64)
points2 = np.array([[x_points2, y, z] for y in y_points2 for z in z_points2], dtype=np.float64)

"""
use the tools and processing gizmos.
"""
# process interpolation/differentiation of points.
result_force1 = getData(dataset, variable1, time1, temporal_method, spatial_method1, spatial_operator1, points1)
result_force2 = getData(dataset, variable1, time2, temporal_method, spatial_method1, spatial_operator1, points2)
result_pressure1 = getData(dataset, variable2, time1, temporal_method, spatial_method2, spatial_operator2, points1)
result_pressure2 = getData(dataset, variable2, time2, temporal_method, spatial_method2, spatial_operator2, points2)
```

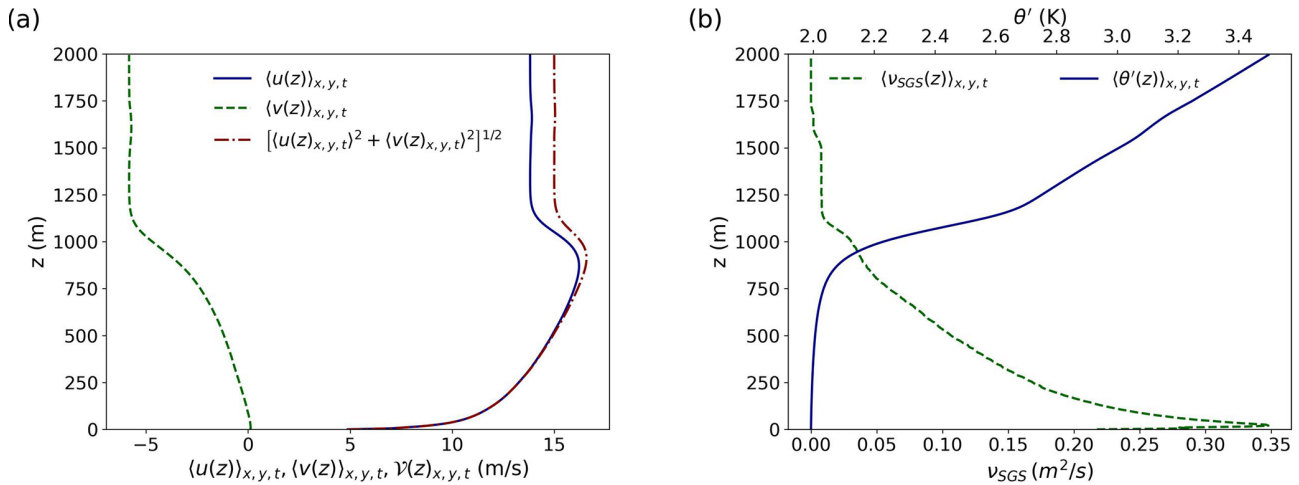**Figure 9.** Python code snippet used to obtain the data to generate Fig. 8.

Figure 10. Vertical profiles of horizontal- and time-averaged **(a)** velocities $\langle u(z)\rangle_{x,y,t}$, $\langle v(z)\rangle_{x,y,t}$ and velocity magnitude $\mathcal{V}(z)_{x,y,t} = \left[\langle u(z)_{x,y,t}\rangle^2 + \langle v(z)_{x,y,t}\rangle^2\right]^{1/2}$, **(b**, bottom axis) subgrid-scale eddy viscosity $\langle \nu_{SGS}(z)\rangle_{x,y,t}$ used in the LES as a result of the Lagrangian scale-dependent dynamic model, and **(b**, top axis) potential temperature deviation $\langle \theta'(z)\rangle_{x,y,t}$ (i.e., the deviations from a reference temperature $\theta_0 = 263.5$ K).

```python
"""
initialize getData parameters (except time and points)
"""
variable, temporal_method, spatial_method, spatial_operator = 'velocity', 'pchip', 'lag8', 'field'


"""
initialize getData times and points
"""
time1, time2, time3, time4, nx, ny, nz = 900, 1800, 2700, 3598.5, 200, 100, 100

x_start, x_end, y_start, y_end, z_start, z_end = 0, 10584, 0, 3780, 2.5, 1997.5
x_points, y_points, z_points = np.linspace(x_start, x_end, nx, dtype=np.float64), np.linspace(y_start, y_end, ny, dtype=np.float64), np.linspace(z_start, z_end, nz, dtype=np.float64)
points = np.array([[x, y, z] for x in x_points for y in y_points for z in z_points], dtype=np.float64)


"""
use the tools and processing gizmos.
"""
# Process interpolation/differentiation of points
result1 = getData(dataset, variable, time1, temporal_method, spatial_method, spatial_operator, points)
result2 = getData(dataset, variable, time2, temporal_method, spatial_method, spatial_operator, points)
result3 = getData(dataset, variable, time3, temporal_method, spatial_method, spatial_operator, points)
result4 = getData(dataset, variable, time4, temporal_method, spatial_method, spatial_operator, points)

# Calculate the average
average_result = [
    (v1 + v2 + v3 + v4) / 4
    for v1, v2, v3, v4 in zip(result1, result2, result3, result4)
]
```

Figure 11. Python code snippet used to obtain the data to generate vertical profiles of $\langle u(z)\rangle_{x,y,t}$: for the 100 heights z between $z = 2.5$ m and $z = 1997.5$ m, we query data on a regular mesh (not necessarily coinciding with stored grid points). For statistical convergence, we average over four times covering the entire hour ($t = 900$ s, 1800 s, 2700 s, 3598.5 s).

approach would require users to identify specific files, understand naming conventions, and handle formatting, posing a barrier to seamless integration with flow field queries. To maintain consistency and usability across the platform, we adopt a similar virtual sensor data access paradigm used for the flow field data. Two dedicated query functions are developed: `getTurbineData(...)` for turbine-level quantities and `getBladeData(...)` for blade-resolved data. For `getTurbineData (...)`, users specify the turbine number (ranging from 1 to 60) and desired time instances. For `getBladeData (...)`, both turbine number and blade number (ranging from 1 to 3) need to be specified, along with an array of actuator point indices (ranging from 1 to 100) and times (ranging from 1 to 3600 s) at which the

data are requested. Linear interpolation in time is supported to provide values between stored simulation time steps.

As an example, Fig. 12 presents the time series of total wind farm power output (panel a) and of rows 1 and 9 of six turbines (panel b). The code snippet specifying the `getTurbineData(...)` call is shown in Fig. 13. Similar calls can be made to extract any of the turbine-specific variables listed in Table 4.

Next, we illustrate the use of `getBladeData(...)` in Fig. 14, which shows (a) the time evolution of the lift and drag coefficients and (b) the lift coefficient as a function of blade angle. The blade angle is computed as $\zeta(t) = \arctan[z(t) - z_h)/(x(t) - x_T)]$ over a 60 s period. The results shown are for a particular turbine and blade (Tur-
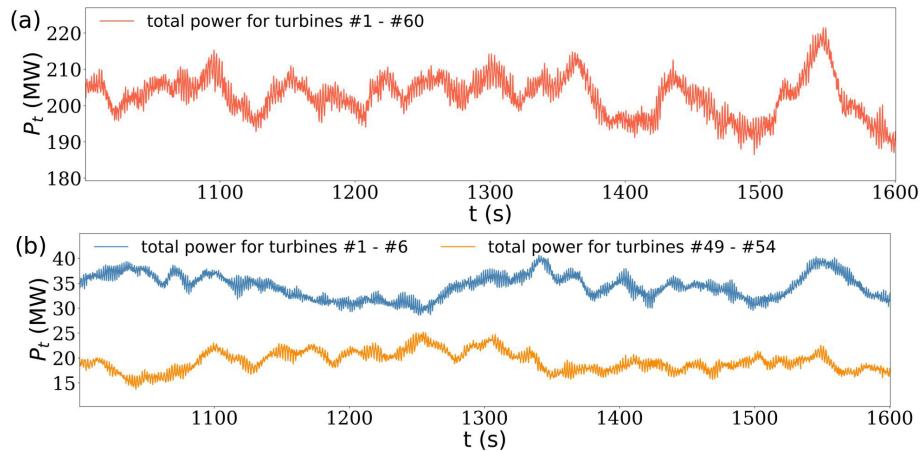
Figure 12. Time evolution of power from turbines during the 10 min time interval, i.e., $t \in [1000.33, 1600.33]$ s. Panel **(a)** shows the total power from the entire wind farm, while panel **(b)** shows the power for the turbines in Row 1 (i.e., turbines #1–#6) and in Row 9 (i.e., turbines #49–#54).

```python
"""
initialize getTurbineData parameters
"""
turbines = list(range(1, 61))
turbine_variable = 'power'

"""
initialize time array, below shows the one from 1000.33 s to 1600.33 s and the time interval is 0.025s
"""
time_start, time_end, dt = 1000.33, 1600.33, 0.025
ntime = int((time_end - time_start) / dt)
turbine_times = np.linspace(time_start, time_end, ntime, dtype=np.float64)

"""
use the tools and processing gizmos.
"""
# process turbine data.
turbine result = getTurbineData(dataset, turbines, turbine variable, turbine times)
```

Figure 13. Python code snippet illustrating the use of the function `getTurbineData(...)` as part of a loop over all turbines in the wind farm and subsequent summation to evaluate time series of total power used to generate Fig. 12a.

bine #28 in the central portion of the wind farm and Blade #3 – the latter being an arbitrary choice, of course). The Python code snippet shown in Fig. 15 illustrates the use of `getBladeData(...)`, with the queried data plotted directly as a time series within the same script. Using a similar approach, variable data can be extracted along turbine blades and further processed to compute higher-order statistics. Figure 16 shows axial force, tangential force, and drag and lift coefficients for an upstream turbine (Blade #1 of Turbine #1) and a downstream turbine (Blade #1 of Turbine #60) at a specific time of $t = 1500$ s. Any of the variables listed in Table 5 can be similarly queried (also in MATLAB).

6 Conclusions

In this paper, we have introduced JHTDB-wind, hosting datasets from high-fidelity LES simulations of wind farms. We extend the standard "virtual sensors" data access methods (Li et al., 2008; Yu et al., 2012; Graham et al., 2016) that have been successfully used for democratizing access to more fundamental turbulence datasets. Besides velocity, pressure, potential temperature, and SGS eddy-viscosity fields, JHTDB-wind adds 4D space–time data on aerodynamic turbine force distributions as seen by the flow and the time series of turbine- and actuator-line-specific aerodynamic data along each of the turbine blades, modeled using ALM. We explain the simulation details, provide background on the numerical method and flow parameters, and provide detailed examples and explanations of the user-friendly data access methodolo-

(a)

(b)

Figure 14. (a) Time evolution of lift and drag coefficients on an ALM point 80 % along the span of Blade #3 for Turbine # 28. **(b)** Polar plot of lift coefficient for that point as a function of blade angle along its rotation. For this turbine, the rotational speed is fixed at $\Omega = 1.09\,\mathrm{rad\,s^{-1}}$ (as obtained from `getTurbineData(...)`), corresponding to approximately 10.5 revolutions during a 60 s period.

```python
"""
initialize getBladeData parameters
"""
turbines, blades, blade_variable1, blade_variable2 = [28], [3], 'Cl', 'Cd'

"""
initialize time array, below shows the one from 1000.33 s to 1600.33 s and the time interval is 0.025s
"""
time_start, time_end, dt = 1000.33, 1600.33, 0.025
ntime = int((time_end - time_start) / dt)
blade_times = np.linspace(time_start, time_end, ntime, dtype=np.float64)
blade_actuator_points = [80]

"""
use the tools and processing gizmos.
"""
# process blade data.
blade_result1 = getBladeData(dataset, turbines, blades, blade_variable1, blade_times, blade_actuator_points)
blade_result2 = getBladeData(dataset, turbines, blades, blade_variable2, blade_times, blade_actuator_points)
```

Figure 15. Python code snippet used to obtain the data to generate Fig. 14.

Figure 16. Distributions of ALM quantities along the turbine blade at a specific time ($t = 1500\,\mathrm{s}$ for Blade #1 of Turbine #1, blue lines; Blade # 1 of Turbine #60, orange lines): **(a)** axial component of the local force (on each $\Delta\ell = 0.615\,\mathrm{m}$ segment) F_{axi}, **(b)** tangential component of the local force (on each $\Delta\ell = 0.615\,\mathrm{m}$ segment) F_{tan}, **(c)** lift coefficient C_{l}, **(d)** drag coefficient C_{d}.

gies. It is hoped that these data will provide useful insights about the complex fluid dynamic processes occurring in wind farms.

We realize that in generating a dataset for a representative conventionally neutral boundary layer case, with a relatively large wind farm with 60 turbines, many other choices could have been made (flow parameters, turbine model and control scheme, usage of a particular LES numerical code, numerical resolution, and so on). We anticipate that different members of the community would have made different choices, and we look forward to conversations about how to further improve such datasets. We believe, however, that the case selected is representative of CNBL wind farm dynamics that have been studied by many others before, with a well-tested numerical code. Hence, the authors hope that the data can be of some use and interest to researchers in wind energy.

As a final note, we have additionally prepared a second dataset for JHTDB-wind featuring an eight-turbine wind farm over a full diurnal cycle, capturing both strongly stable and unstable atmospheric boundary layer regimes at different times of the day and night (Xiao et al., 2025).

Author contributions. XZ performed the simulations, generated the majority of the data, and assisted in document and figure preparation and detailed proof-reading. SX performed the majority of the data transformation into Zarr and Parquet formats, worked on testing data access methods, and generated many of the figures. GN developed the thermal stratification and initialization methods in the LES code. LAMT developed and implemented the generalized ALM method in the LES code. MS and HY developed the giverny backend software and Python/MATLAB data access codes. GL directed the SciServer and Zarr format optimization. AS designed the storage architecture. DG participated in data interpretation and analysis and edited the article. CM participated in simulation and database design, data interpretation and analysis, and document preparation and proof-reading.

References

Abkar, M. and Porté-Agel, F.: The effect of free-atmosphere stratification on boundary-layer flow and power output from very large wind farms, Energies, 6, 2338–2361, https://doi.org/10.3390/en6052338, 2013.

Abkar, M. and Porté-Agel, F.: Mean and turbulent kinetic energy budgets inside and above very large wind farms under conventionally-neutral condition, Renewable Energy, 70, 142–152, https://doi.org/10.1016/j.renene.2014.03.050, 2014.

Aitken, M. L., Kosović, B., Mirocha, J. D., and Lundquist, J. K.: Large eddy simulation of wind turbine wake dynamics in the stable boundary layer using the Weather Research and Forecasting Model, J. Renewable Sustainable Energy, 6, https://doi.org/10.1063/1.4885111, 2014.

Aiyer, A. K., Deike, L., and Mueller, M. E.: A dynamic wall modeling approach for large eddy simulation of offshore wind farms in realistic oceanic conditions, J. Renewable Sustainable Energy, 16, https://doi.org/10.1063/5.0159019, 2024.

Albertson, J. D.: Large eddy simulation of land-atmosphere interaction, Ph.D. thesis, University of California, Davis, 1996.

Albertson, J. D. and Parlange, M. B.: Surface length scales and shear stress: Implications for land-atmosphere interaction over complex terrain, Water Resour. Res., 35, 2121–2132, https://doi.org/10.1029/1999WR900094, 1999.

Alexakis, A., Marino, R., Mininni, P. D., van Kan, A., Foldes, R., and Feraco, F.: Large-scale self-organization in dry turbulent atmospheres, Science, 383, 1005–1009, https://doi.org/10.1126/science.adg8269, 2024.

Allaerts, D. and Meyers, J.: Large eddy simulation of a large wind-turbine array in a conventionally neutral atmospheric boundary layer, Phys. Fluids, 27, https://doi.org/10.1063/1.4922339, 2015.

Allaerts, D. and Meyers, J.: Boundary-layer development and gravity waves in conventionally neutral wind farms, J. Fluid Mech., 814, 95–130, https://doi.org/10.1017/jfm.2017.11, 2017.

Ayala, M., Sadek, Z., Ferčák, O., Cal, R. B., Gayme, D. F., and Meneveau, C.: A moving surface drag model for LES of wind over waves, Boundary Layer Meteorol., 190, 39, https://doi.org/10.1007/s10546-024-00884-8, 2024.

Bodini, N., Optis, M., Redfern, S., Rosencrans, D., Rybchuk, A., Lundquist, J. K., Pronk, V., Castagneri, S., Purkayastha, A., Draxl, C., Krishnamurthy, R., Young, E., Roberts, B., Rosenlieb, E., and Musial, W.: The 2023 National Offshore Wind data set (NOW-23), Earth Syst. Sci. Data, 16, 1965–2006, https://doi.org/10.5194/essd-16-1965-2024, 2024.

Bou-Zeid, E., Meneveau, C., and Parlange, M. B.: Large-eddy simulation of neutral atmospheric boundary layer flow over heterogeneous surfaces: Blending height and effective surface roughness, Water Resour. Res., 40, https://doi.org/10.1029/2003WR002475, 2004.

Bou-Zeid, E., Meneveau, C., and Parlange, M.: A scale-dependent Lagrangian dynamic model for large eddy simulation of complex turbulent flows, Phys. Fluids, 17, https://doi.org/10.1063/1.1839152, 2005.

Calaf, M., Meneveau, C., and Meyers, J.: Large eddy simulation study of fully developed wind-turbine array boundary layers, Phys. Fluids, 22, https://doi.org/10.1063/1.3291077, 2010.

Chatelain, P., Backaert, S., Winckelmans, G., and Kern, S.: Large eddy simulation of wind turbine wakes, Flow Turbul. Combust., 91, 587–605, https://doi.org/10.1007/s10494-013-9474-8, 2013.

Chung, W. T., Jung, K. S., Chen, J. H., and Ihme, M.: BLAST-Net: A call for community-involved big data in combustion machine learning, Appl. Energy Combust. Sci., 12, 100 087, https://doi.org/10.1016/j.jaecs.2022.100087, 2022.

Churchfield, M., Lee, S., Moriarty, P., Martinez, L., Leonardi, S., Vijayakumar, G., and Brasseur, J.: A large-eddy simulation of wind-plant aerodynamics, in: 50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition, 537, https://doi.org/10.2514/6.2012-537, 2012.

Duraisamy, K., Iaccarino, G., and Xiao, H.: Turbulence modeling in the age of data, Annu. Rev. Fluid Mech., 51, 357–377, https://doi.org/10.1146/annurev-fluid-010518-040547, 2019.

Durran, D. R. and Klemp, J. B.: A compressible model for the simulation of moist mountain waves, Mon. Weather Rev., 111, 2341–2361, https://doi.org/10.1175/1520-0493(1983)111<2341:ACMFTS>2.0.CO;2, 1983.

Gadde, S. N. and Stevens, R. J.: Interaction between low-level jets and wind farms in a stable atmospheric boundary layer, Phys. Rev. Fluids, 6, 014603, https://doi.org/10.1103/PhysRevFluids.6.014603, 2021.

Gebraad, P. M., Teeuwisse, F. W., Van Wingerden, J., Fleming, P. A., Ruben, S. D., Marden, J. R., and Pao, L. Y.: Wind plant power optimization through yaw control using a parametric model for wake effects – a CFD simulation study, Wind Energy, 19, 95–114, https://doi.org/10.1002/we.1822, 2016.

Gharaati, M., Xiao, S., Wei, N. J., Martínez-Tossas, L. A., Dabiri, J. O., and Yang, D.: Large-eddy simulation of helical- and straight-bladed vertical-axis wind turbines in boundary layer turbulence, J. Renewable Sustainable Energy, 14, https://doi.org/10.1063/5.0100169, 2022.

Gharaati, M., Xiao, S., Martínez-Tossas, L. A., Araya, D. B., and Yang, D.: Large-eddy simulations of turbulent wake flows behind helical-and straight-bladed vertical axis wind turbines rotating at low tip speed ratios, Phys. Rev. Fluids, 9, 074603, https://doi.org/10.1103/PhysRevFluids.9.074603, 2024.

Graham, J., Kanov, K., Yang, X., Lee, M., Malaya, N., Lalescu, C., Burns, R., Eyink, G., Szalay, A., Moser, R., and Meneveau, C.: A web services accessible database of turbulent channel flow and its use for testing a new integral wall model for LES, J. Turbul., 17, 181–215, https://doi.org/10.1080/14685248.2015.1088656, 2016.

Howland, M. F., Bossuyt, J., Martínez-Tossas, L. A., Meyers, J., and Meneveau, C.: Wake structure in actuator disk models of wind turbines in yaw under uniform inflow conditions, J. Renewable Sustainable Energy, 8, https://doi.org/10.1063/1.4955091, 2016.

Jha, P. K., Churchfield, M. J., Moriarty, P. J., and Schmitz, S.: Guidelines for volume force distributions within actuator line modeling of wind turbines on large-eddy simulation-type grids, J. Sol. Energy Eng., 136, 031003, https://doi.org/10.1115/1.4026252, 2014.

Jonkman, J., Butterfield, S., Musial, W., and Scott, G.: Definition of a 5-MW reference wind turbine for offshore system development, Technical Report NREL/TP-500-38060, National Renewable Energy Laboratory (NREL), Golden, CO, USA, 2009.

Kosović, B. and Curry, J. A.: A large eddy simulation study of a quasi-steady, stably stratified atmospheric boundary layer, J. Atmos. Sci., 57, 1052–1068, https://doi.org/10.1175/1520-0469(2000)057<1052:ALESSO>2.0.CO;2, 2000.

Kumar, V., Kleissl, J., Meneveau, C., and Parlange, M. B.: Large-eddy simulation of a diurnal cycle of the atmospheric boundary layer: Atmospheric stability and scaling issues, Water Resour. Res., 42, https://doi.org/10.1029/2005WR004651, 2006.

Kusiak, A.: Renewables: Share data on wind energy, Nature, 529, 19–21, 2016.

Li, C., Liu, L., Lu, X., and Stevens, R. J.: Analytical model of fully developed wind farms in conventionally neutral atmospheric boundary layers, J. Fluid Mech., 948, A43, https://doi.org/10.1017/jfm.2022.732, 2022.

Li, Y., Perlman, E., Wan, M., Yang, Y., Meneveau, C., Burns, R., Chen, S., Szalay, A., and Eyink, G.: A public turbulence database cluster and applications to study Lagrangian evolution of velocity increments in turbulence, J. Turbul., 9, 31, https://doi.org/10.1080/14685240802376389, 2008.

Lilly, D.: The representation of small-scale turbulence in numerical simulation experiments, Technical report, National Center for Atmospheric Research (NCAR), 1966.

Liu, L. and Stevens, R. J.: Vertical structure of conventionally neutral atmospheric boundary layers, P. Natl. Acad. Sci. USA, 119, e2119369119, https://doi.org/10.1073/pnas.2119369119, 2022.

Liu, L., Lu, X., and Stevens, R. J.: Geostrophic drag law in conventionally neutral atmospheric boundary layer: simplified parametrization and numerical validation, Boundary Layer Meteorol., 190, 37, https://doi.org/10.1007/s10546-024-00878-6, 2024.

Martinez, L. A.: Large eddy simulations and theoretical analysis of wind turbine aerodynamics using an actuator line model, Ph.D. thesis, Johns Hopkins University, 2017.

Martínez-Tossas, L., Churchfield, M., and Meneveau, C.: Optimal smoothing length scale for actuator line models of wind turbine blades based on Gaussian body force distribution, Wind Energy, 20, 1083–1096, https://doi.org/10.1002/we.2081, 2017.

Martínez-Tossas, L. A. and Meneveau, C.: Filtered lifting line theory and application to the actuator line model, J. Fluid Mech., 863, 269–292, https://doi.org/10.1017/jfm.2018.994, 2019.

Martínez-Tossas, L. A., Churchfield, M. J., and Meneveau, C.: Large eddy simulation of wind turbine wakes: detailed comparisons of two codes focusing on effects of numerics and subgrid modeling, in: J. Phys.: Conference Series, IOP Publishing, 625, 012024, https://doi.org/10.1088/1742-6596/625/1/012024, 2015.

Martínez-Tossas, L. A., Sakievich, P., Churchfield, M. J., and Meneveau, C.: Generalized filtered lifting line theory for arbitrary chord lengths and application to wind turbine blades, Wind Energy, 27, 101–106, https://doi.org/10.1002/we.2872, 2024.

McWilliams, J. C., Weiss, J. B., and Yavneh, I.: Anisotropy and coherent vortex structures in planetary turbulence, Science, 264, 410–413, https://doi.org/10.1126/science.264.5157.410, 1994.

Meyers, J. and Meneveau, C.: Optimal turbine spacing in fully developed wind farm boundary layers, Wind Energy, 15, 305–317, https://doi.org/10.1002/we.469, 2012.

Miles, A., jakirkham, Bussonnier, M., Moore, J., Papadopoulos Orfanos, D., Bourbeau, J., Fulton, A., Lee, G., Patel, Z., Bennett, D., Rocklin, M., AWA BRANDON AWA, Chopra, S., Abernathey, R., Kristensen, M. R. B., Sales de Andrade, E., Durant, M., Schut, V., Dussin, R., Verma, S., Chaudhary, S., Barnes, C., Hamman, J., Nunez-Iglesias, J., Williams, B., Mohar, B., Noyes, C., and Bolarinwa, E.: zarr-developers/zarr-python: v2.15.0, Zenodo [code], https://doi.org/10.5281/zenodo.8039103, 2023.

Monin, A. S. and Obukhov, A. M.: Basic laws of turbulent mixing in the surface layer of the atmosphere, Contrib. Geophys. Inst. Acad. Sci. USSR, 151, e187, 1954.

Munters, W., Meneveau, C., and Meyers, J.: Shifted periodic boundary conditions for simulations of wall-bounded turbulent flows, Phys. Fluids, 28, https://doi.org/10.1063/1.4941912, 2016.

Narasimhan, G., Gayme, D. F., and Meneveau, C.: Effects of wind veer on a yawed wind turbine wake in atmospheric boundary layer flow, Phys. Rev. Fluids, 7, 114609, https://doi.org/10.1103/PhysRevFluids.7.114609, 2022.

Narasimhan, G., Gayme, D. F., and Meneveau, C.: Analytical wake modeling in atmospheric boundary layers: accounting for wind veer and thermal stratification, J. Phys.: Conference Series, IOP Publishing, 2767, 092018, https://doi.org/10.1088/1742-6596/2767/9/092018, 2024a.

Narasimhan, G., Gayme, D. F., and Meneveau, C.: Analytical model coupling Ekman and surface layer structure in atmospheric boundary layer flows, Boundary Layer Meteorol., 190, 16, https://doi.org/10.1007/s10546-024-00859-9, 2024b.

Narasimhan, G., Gayme, D. F., and Meneveau, C.: An extended analytical wake model and applications to yawed wind tur-

bines in atmospheric boundary layers with different levels of stratification and veer, J. Renewable Sustainable Energy, 17, https://doi.org/10.1063/5.0251305, 2025.

NCAR: HPE SGI ICE XA – Cheyenne, NCAR, https://doi.org/10.5065/D6RX99HX, 2025.

Perlman, E., Burns, R., Li, Y., and Meneveau, C.: Data exploration of turbulence simulations using a database cluster, in: Proceedings of the 2007 ACM/IEEE Conference on Supercomputing, 1–11, https://doi.org/10.1145/1362622.1362654, 2007.

Porté-Agel, F., Meneveau, C., and Parlange, M. B.: A scale-dependent dynamic model for large-eddy simulation: application to a neutral atmospheric boundary layer, J. Fluid Mech., 415, 261–284, https://doi.org/10.1017/S0022112000008776, 2000.

Sescu, A. and Meneveau, C.: A control algorithm for statistically stationary large-eddy simulations of thermally stratified boundary layers, Q. J. R. Meteorolog. Soc., 140, 2017–2022, https://doi.org/10.1002/qj.2266, 2014.

Shapiro, C. R., Gayme, D. F., and Meneveau, C.: Modelling yawed wind turbine wakes: a lifting line approach, J. Fluid Mech., 841, R1, https://doi.org/10.1017/jfm.2018.75, 2018.

Shapiro, C. R., Gayme, D. F., and Meneveau, C.: Generation and decay of counter-rotating vortices downstream of yawed wind turbines in the atmospheric boundary layer, J. Fluid Mech., 903, R2, https://doi.org/10.1017/jfm.2020.717, 2020.

Smagorinsky, J.: General circulation experiments with the primitive equations: I. The basic experiment, Mon. Weather Rev., 91, 99–164, https://doi.org/10.1175/1520-0493(1963)091<0099:GCEWTP>2.3.CO;2, 1963.

Sørensen, J. N. and Shen, W. Z.: Numerical modeling of wind turbine wakes, J. Fluids Eng., 124, 393–399, https://doi.org/10.1115/1.1471361, 2002.

Stevens, R. J. and Meneveau, C.: Flow structure and turbulence in wind farms, Annu. Rev. Fluid Mech., 49, 311–339, https://doi.org/10.1146/annurev-fluid-010816-060206, 2017.

Stevens, R. J., Graham, J., and Meneveau, C.: A concurrent precursor inflow method for large eddy simulations and applications to finite length wind farms, Renewable Energy, 68, 46–50, https://doi.org/10.1016/j.renene.2014.01.024, 2014.

Stevens, R. J., Martínez-Tossas, L. A., and Meneveau, C.: Comparison of wind farm large eddy simulations using actuator disk and actuator line models with wind tunnel experiments, Renewable Energy, 116, 470–478, https://doi.org/10.1016/j.renene.2017.08.072, 2018.

Troldborg, N.: Actuator line modeling of wind turbine wakes, Ph.D. thesis, Technical University of Denmark, 2009.

Xiao, S., Zhu, X., Narasimhan, G., Gayme, D. F., and Meneveau, C.: Wind farm dynamics over a diurnal cycle: analysis of a comprehensive large eddy simulation, web-services accessible dataset, arXiv [preprint], https://doi.org/10.48550/arXiv.2510.05005, 2025.

Yang, D., Meneveau, C., and Shen, L.: Large-eddy simulation of offshore wind farm, Phys. Fluids, 26, https://doi.org/10.1063/1.4863096, 2014.

Yang, X., Milliren, C., Kistner, M., Hogg, C., Marr, J., Shen, L., and Sotiropoulos, F.: High-fidelity simulations and field measurements for characterizing wind fields in a utility-scale wind farm, Appl. Energy, 281, 116115, https://doi.org/10.1016/j.apenergy.2020.116115, 2021.

Yu, H., Kanov, K., Perlman, E., Graham, J., Frederix, E., Burns, R., Szalay, A., Eyink, G., and Meneveau, C.: Studying Lagrangian dynamics of turbulence using on-demand fluid particle tracking in a public turbulence database, J. Turbul., 13, 12, https://doi.org/10.1080/14685248.2012.674643, 2012.

Zhang, C., Duan, L., and Choudhari, M. M.: Direct numerical simulation database for supersonic and hypersonic turbulent boundary layers, AIAA Journal, 56, 4297–4311, https://doi.org/10.2514/1.J057296, 2018.

Zhang, F., Yang, X., and He, G.: Multiscale analysis of a very long wind turbine wake in an atmospheric boundary layer, Phys. Rev. Fluids, 8, 104605, https://doi.org/10.1103/PhysRevFluids.8.104605, 2023.

Zhu, X., Xiao, S., Narasimhan, G., Martinez-Tossas, L. A., Schnaubelt, M., Lemson, G., Szalay, A., Gayme, D. F., and Meneveau, C.: Large wind farm under 1-hour conventionally neutral atmospheric conditions Johns Hopkins Turbulence Databases – Wind [data set], https://doi.org/10.26144/D8ES-FC15, 2025.

Zilitinkevich, S., Baklanov, A., Rost, J., Smedman, A.-s., Lykosov, V., and Calanca, P.: Diagnostic and prognostic equations for the depth of the stably stratified Ekman boundary layer, Q. J. R. Meteorolog. Soc., 128, 25–46, https://doi.org/10.1256/00359000260498770, 2002.