A baseline for ensemble-based, time-resolved inflow reconstruction for a single turbine using large-eddy simulations and latent diffusion models

A Rybchuk, L A Martínez-Tossas, N Hamilton, P Doubrawa, G Vijayakumar, M Hassanaly, M B Kuhn, D S Zalkind

National Renewable Energy Laboratory, 15013 Denver W Pkwy, Golden, CO 80401, USA

E-mail: alex.rybchuk@nrel.gov

Abstract. We are interested in reconstructing winds flowing through a turbine on a second-by-second basis over a 10 min window. Previously, we developed a machine learning algorithm that takes in a snapshot of wind speed measurements and generates ensembles of three-dimensional wind field estimates. Here, we use these estimates as initial conditions in large-eddy simulations and reconstruct atmospheric and turbine response dynamic quantities in a synthetic field campaign. In doing so, we establish a baseline for model validation that future time-aware data assimilation techniques will be compared to. In turbine-free case studies, ground truth wind speeds consistently fall within our estimated wind speed distribution for the first 100 s after the simulation start. In simulations with turbines, the wind estimates show a small bias of 0.10 m s^{-1} and good correlation of 0.80 during the first 100 s. During this window, our estimates of the Blade 1 bending moment and generator power typically span the ground truth, with the estimate of the former performing better overall. In summary, this approach shows promise as a stand-alone technique for reconstructing real-world inflow and turbine dynamics in 1–2 min windows and as a foundation for future time-aware data assimilation techniques.

1. Introduction

The Rotor Aerodynamics Aeroelastics and Wake (RAAW) field campaign aims to gather the most detailed atmospheric and structural measurements of a single land-based, utility-scale wind turbine conducted to date. While turbine validation campaigns often assess time-averaged quantities (e.g., [1]), RAAW aims to validate a computational model of a turbine against observed second-by-second dynamics over a 10 min window. To accomplish this, first the time history of winds at the unmeasured inflow plane immediately upwind of the rotor disk must be accurately reconstructed using a combination of measurements and atmospheric modeling. To this end, the field campaign has deployed a network of atmospheric instruments upwind of the turbine: a meteorological mast, a nacelle-mounted scanning lidar, and a hub-mounted "spinner" lidar.

In previous work [2], we laid the foundation for an ensemble-based flow reconstruction approach using a machine learning algorithm known as a latent diffusion model (LDM) [3]. Specifically, we trained an LDM to read in a snapshot of synthetic observations from large-eddy simulation (LES) output and then generate an ensemble of plausible three-dimensional wind field estimates. While an LDM-generated estimate produces wind fields for only one moment

in time, it can be plugged into an LES code as a turbulent initial condition that reflects the estimated state in the vicinity of the atmospheric observation network. This initial condition enables us to pursue two strategies for time-history reconstruction. In the first category, good initial conditions encourage the success of time-aware model-observation fusion techniques, such as nudging [4] and the ensemble Kalman filter [5]. In the second category, good initial conditions enable us to take a "guess-and-check" approach to time-history reconstruction without the need to implement a time-aware data fusion technique—i.e., given several initial conditions, we can simulate many time histories, and then after the fact, we can select the initial conditions that produce the greatest similarity in target dynamics by comparing between inflow measurements in the ground truth and state estimation simulations.

Here, we explore the question: With a guess-and-check approach enabled by the LDM, how accurately can we reconstruct the history of inflow winds and turbine dynamics in a synthetic LES study? We treat this approach as a baseline against which we will compare future time-aware reconstruction techniques. To that end, we simulate several observation periods of a synthetic field campaign using LES during a period that the LDM has not been trained on. Using ensembles of initial conditions to drive LES runs, we attempt to reconstruct several time-varying quantities in five case studies without a turbine and five case studies with a turbine. In section 2, we describe the RAAW field campaign, the LES approximation of the field campaign, the flow reconstruction problem studied here, and our flow reconstruction strategy in greater detail. In section 3, we compare our estimates of dynamics to the ground truth in simulations with and without the wind turbine. In section 4, we present conclusions and discuss next possible steps.

2. Methods

2.1. RAAW field campaign

The RAAW field campaign seeks to characterize the behavior of a 2.8 MW GE wind turbine with a rotor diameter of $D = 127$ m. The campaign has deployed a network of atmospheric sensors to measure flow fields in the vicinity of the turbine (figure 2b,e). A meteorological mast is located approximately $3D$ upwind of the turbine along the predominant wind direction for the site; the mast measures u, v, and w between 30 m and 185 m above ground. Two scanning lidars sit atop the turbine nacelle, one pointed upwind and one pointed downwind, measuring the line-of-sight velocity up to a range of approximately $8D$. We work with only the upwind lidar here. Finally, a continuous wave "spinner" lidar measures line-of-sight velocity over an area the size of the rotor disk at approximately $1D$ upwind of the turbine.

2.2. Synthetic field campaign

We simulate one set of atmospheric conditions pertinent to the RAAW field campaign using the LES code AMR-Wind [6]. Following [2], we simulate thermally neutral atmospheric boundary layers forced by $(U, V) = (10, 0)$ m s^{-1} geostrophic winds, no surface heating, and Coriolis forcing that would be appropriate at a latitude of 90°. Simulations without turbines use a time step of 0.5 s, and simulations with turbines have a time step of 0.05 s. Our domain is sized $(L_x, L_y, L_z) = (1920, 1920, 960)$ m, with $(n_x, n_y, n_z) = (128, 128, 64)$ grid cells and 15 m grid spacing in all directions, a spacing that roughly matches the scanning lidar range gate size in the real-world field campaign. We note that this is a relatively small domain, and as a result, there is potential for autocorrelation to occur in cycles of roughly $T = 1920$ m $/10$ m s$^{-1} = 192$ s, although we will demonstrate that this autocorrelation does not appear in practice. The simulations have a capping inversion that sits near a height of 480 m. We approximate the RAAW turbine as an open-source NREL 2.8 MW turbine [7], with an 86.5 m hub height and 127 m rotor diameter, and we position it at $(x, y) = (1020, 960)$ m. In our turbine simulations, we refine the grid to 3.75 m resolution in the vicinity of the turbine.

We approximate the real-world atmospheric measurements in our LES. Our meteorological mast sits 360 m upwind of the turbine and measures all three velocity components (u, v, w) between 0 m and 180 m above ground. Our scanning lidar directly measures the u velocity (instead of line-of-sight velocity) in a circular sector extending 1005 m upwind with a half-angle of 18 degrees. The virtual spinner lidar also measures the u velocity (instead of line-of-sight velocity) in a plane 120 m upwind of the turbine, in a circular area that sweeps the rotor disk. We believe the line-of-sight assumption is warranted in this study, as the lidars are aligned with the incoming flow, which minimizes differences between u and the line-of-sight velocity. However, this assumption will more rigorously be checked in future work involving real-world measurements. All three synthetic instruments output instantaneous velocities at every time step of the simulation, and both lidars output instantaneous snapshots (instead of performing sweeping scans). We approximate the velocity measurements from all instruments as being the same as the LES grid cell values.

2.3. Problem definition and case studies

In the RAAW field campaign, we are interested in reconstructing inflow winds (i.e., the upstream winds that will flow into the turbine) and turbine dynamics on a second-by-second scale over a 10 min period. We run 10 case studies in total, five of which use observations that come from atmospheres without a turbine (section 3.1) and five of which use observations that come from atmospheres with a turbine (section 3.2). In the simpler former case, we are only interested in reconstructing time-varying atmospheric dynamics. In the latter, more complicated case, we are interested in reconstructing both atmospheric and turbine dynamics. For both situations, we initialize ground truth simulations 0.5–2.5 days after the end of the training data period of 3.5 days (Sec. 2.4). Each simulation is separated from the others by 0.5 days of simulated time.

After initialization, in the case studies without turbines, we generate ground truth scenarios by simply continuing to run the LES for 10 min. The case studies with turbines require more work to generate the ground truth data. Following common practice, we first run a turbine-free "precursor" simulation with periodic boundary conditions for the x- and y-planes for each case study. We then run a second simulation that includes a turbine. This simulation uses inflow planes from the precursor as the inflow boundary condition and a pressure boundary condition for the outflow planes [8]. While we are interested in flow reconstruction over a 10 min window, we run precursor and turbine simulations for 15 min, taking the synthetic measurement snapshot 5 min into the ground truth turbine simulation, thereby allowing the turbine to equilibrate after a short initial transient phase.

In the context of the entire field campaign, there are many potential quantities of interest that we could attempt to validate our estimates against. Perhaps the most important atmospheric quantity to reconstruct is the time-varying u velocity upstream of the turbine at the location of a measurement, and here we examine u spatially averaged over the spinner lidar disk. Because we have measurements at this location, this would enable us to search for the reconstructions that best match measurements in a real-world field campaign. We look at this quantity in all 10 case studies. Regarding turbine-oriented quantities of interest, we also examine two quantities in the five turbine case studies: generator power production and generator torque.

2.4. Reconstruction approach

The flow reconstruction process here (figure 1) is enabled by a machine learning algorithm known as a latent diffusion model (LDM) [3]. LDMs are a recent deep generative modeling algorithm that excel at producing diverse samples from the probability density function $p(x)$ that characterizes the training data set $\{x_1, x_2, ..., x_N\}$. Given additional information y, LDMs can draw samples from the conditional distribution $p(x|y)$. Here, x represents the full 3D flow field everywhere in an LES simulation, and y represents our synthetic observations. We develop

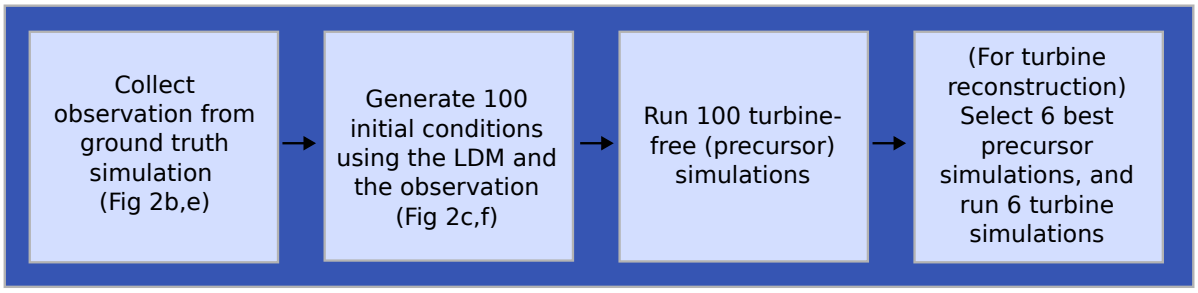

Figure 1. A summary of the flow reconstruction methodology.

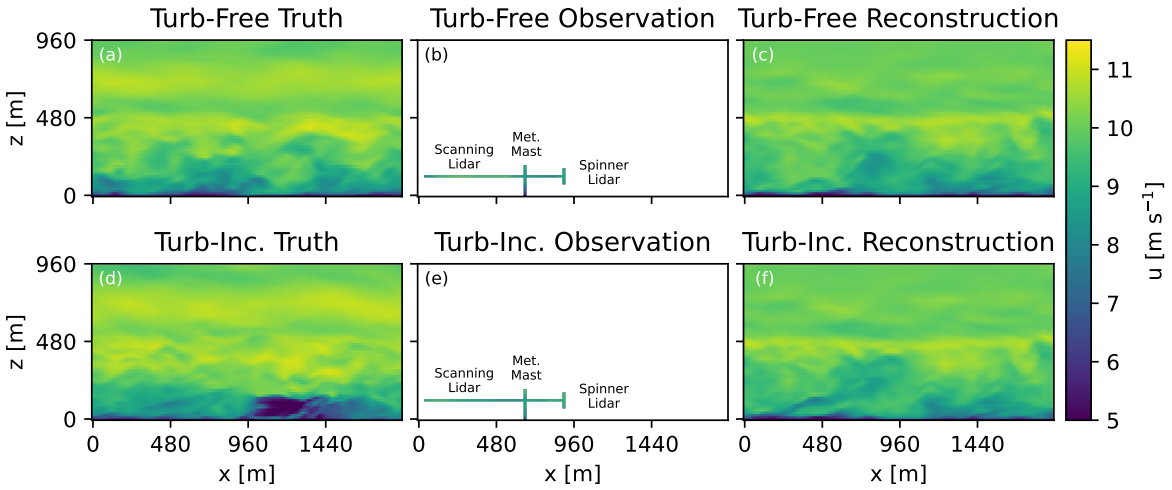

Figure 2. A visual demonstration of the initial condition reconstruction approach in the (a–c) case studies without turbines and (d–f) the case studies with turbines. The panels show an xz-cross-section at the middle of the domain. (a,d) The ground truth initial condition for one turbine-free and one turbine-including case study. (b,e) An observation is generated by masking the ground truth. (c,f) Ensemble members are then generated by feeding the observation into the LDM, and we show one ensemble member. The turbine-including reconstruction (f) omits the wake because the LDM was trained on turbine-free data.

our LDM to reconstruct a snapshot of winds (u, v, w) of the entire 3D LES domain given a snapshot of observations using a mask that matches the RAAW field campaign (figure 2). We train our LDM on turbine-free LES data from one long simulation with the same large-scale forcing as the ground truth simulations, using outputs from every minute between a simulation time of 0.5 and 3.5 days.

We conduct flow reconstruction in several steps. First, given one ground truth observation for each case study, we generate 100 distinct initial conditions. Second, for each generated ensemble member, we run a 10 min simulation without a turbine in it. In section 3.1, we refer to this as a "turbine-free" simulation, and in section 3.2 we refer to this as a "precursor" simulation. Third, for the turbine-focused studies, we additionally run a 10 min turbine simulation. For all reconstructions, we assume that we accurately know the large-scale forcing of the ground truth, so each precursor simulation is also forced with $(U, V) = (10, 0)$ m s^{-1} geostrophic winds and no surface heating.

To estimate the turbine behavior for each turbine case study, we select the six best ensemble members and then run a 10 min turbine simulation using the boundary conditions from their

respective precursor simulations. Thus in section 3.2, we show wind plots with all 100 wind speed estimates from precursor simulations, the six best wind speed estimates, and the six turbine reconstructions from turbine simulations. We opt for six ensemble members, striking a balance between more robustly characterizing uncertainty and carrying out expensive turbine simulations. While there are many possible approaches to define "best ensemble member," here we compare the spatially averaged ground truth and estimated spinner lidar u measurements, a strategy that would be possible in a real-world field campaign. In each case study, we partition the six ensemble members as follows—we select three ensemble members with the smallest spinner lidar u bias (which is perhaps the most intuitive definition of "best") and three ensemble members with the best spinner lidar u correlation (which serves as a proxy for matching turbulence intensity). As we will demonstrate in a moment, estimates are best during the first 100 s, and as such, we choose to calculate these metrics during this window.

3. Results

3.1. Turbine-free wind reconstruction

Before comparing to the ground truth, we first discuss the dynamics of the wind estimates in case studies where observations for the LDM are provided by simulations without turbines (figure 3). Looking across all the case studies, the dynamics of the total estimate ensemble is consistent. For approximately the first 100 s, the ensemble members are grouped together, and afterward the trajectories diverge. This is clearly demonstrated by the ensemble average estimate, which flattens out during this second phase. This performance window of 100 s also demonstrates that the potential autocorrelation that occurs on the order of 192 s because of the small domain size is not an issue in practice. Quantitatively, the difference between the estimated minimum and maximum averaged spinner velocities is approximately 1.25 m s^{-1} during the first 100 s, and approximately 2.0 m s^{-1} afterward. Similarly, the standard deviation between all members is approximately 0.25 m s^{-1} in the first phase, and 0.375 m s^{-1} in the second phase. While this two-phase behavior can be less pronounced when down-selecting to only the dynamics of the six best ensemble members, it is still present. The time correlation value of 100 s has a physical basis. The scanning lidar measures 1005 m upstream, and the hub-height wind speed is approximately 10 m s^{-1}, which means it takes approximately 100 s for winds to traverse the scanning lidar to the spinner lidar. Aside from the two-phase behavior, we also highlight that in some cases the best guesses are initially more tightly coupled (e.g., Case 2) than in others (e.g., Case 1). This behavior suggests that certain flow features are more predictable than others.

In general, the estimates do a good job of matching the ground truth winds in the turbine-free simulations. During the first 100 s, the ground truth falls within the bounds of the 100-member ensemble, and it is almost always spanned in the six-member ensemble. The six-member ensemble mean in particular initially tracks the ground truth well. During the first phase, the three small-bias members never have a bias that exceeds 0.15 m s^{-1}. Similarly, the three best-correlation members have a correlation value of at least $\rho = 0.85$. While the estimates generally perform well, the ground truth infrequently exceeds the value of the 100-member ensemble mean, and across all simulations, the estimates show a positive bias of 0.23 m s^{-1} during the first 100 s. Thus, this bias must be accounted for when reconstructing turbine-free winds, possibly with the use of postprocessing techniques that are commonly applied to ensemble operational weather forecasts [9].

3.2. Turbine-based reconstructions

We next examine flow reconstruction where the observations come from simulations with turbines. In the precursor runs generated using these observations (figure 4 left, center), the estimated winds at the spinner lidar behave similarly to the cases without turbines, except now precursor simulations show a more pronounced bias. The estimates here show the same

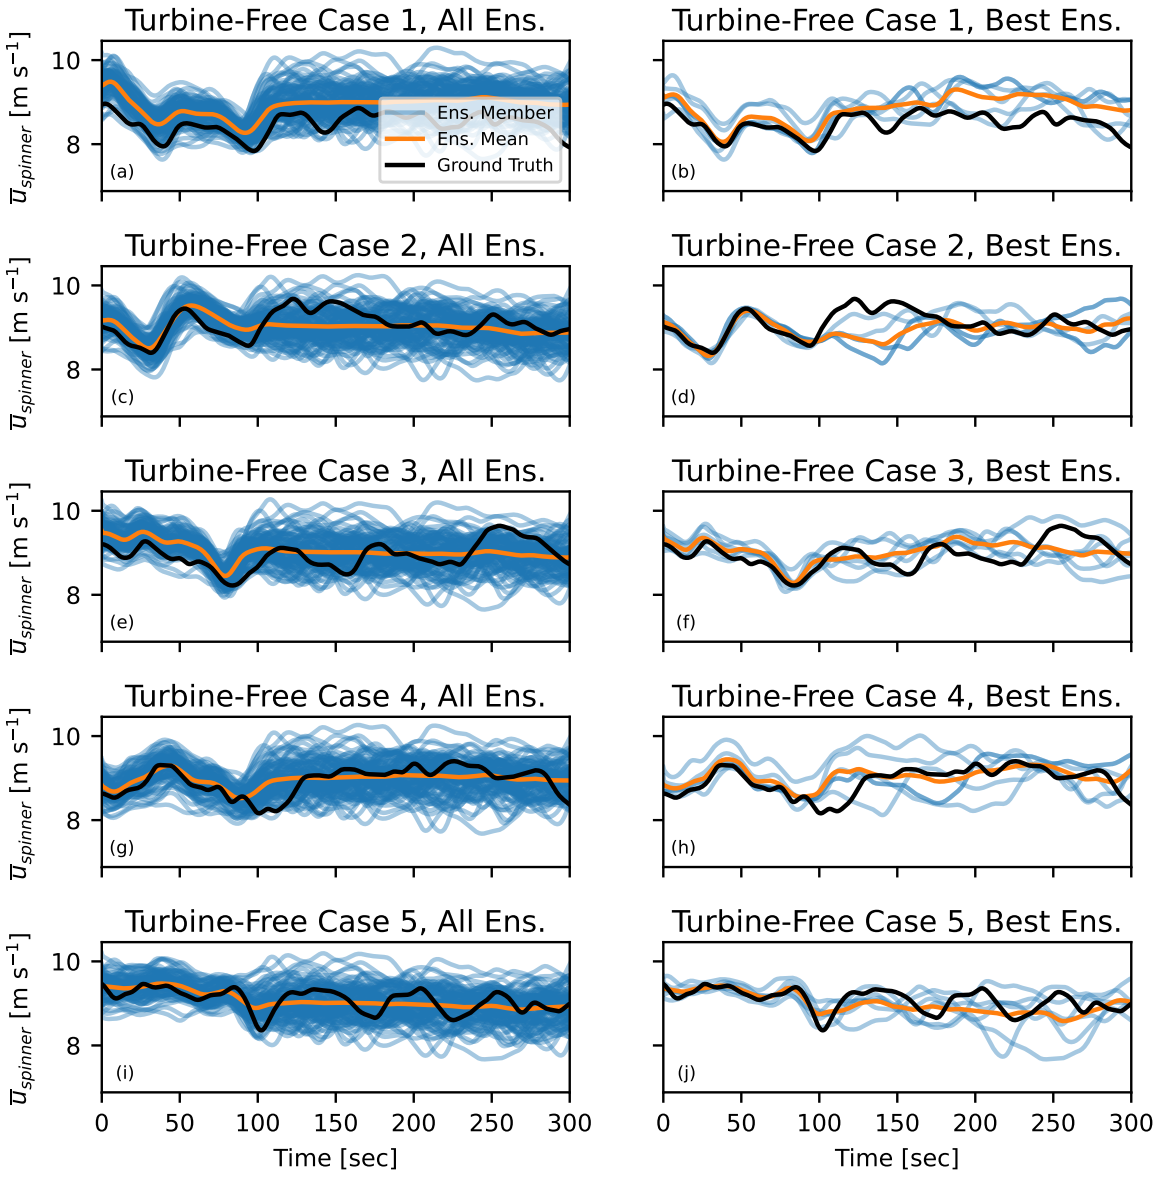

Figure 3. Spatially averaged spinner lidar measurements in the five cases without turbines for the ground truth simulations, (left) 100-member ensemble, and (right) ensemble of six best members. While the ground truth is the same across columns, ensemble means are respectively calculated using 100- and 6-member ensembles.

two-period behavior as before, where estimates initially agree well with one another but then diverge after 100 s. Crucially, the estimates in precursor simulations are significantly more biased relative to the ground truth here than in the turbine-free scenario—full estimate ensembles are now positively biased 0.47 m s^{-1}, nearly double the previous bias. This is expected, as the LDM was trained for a turbine-free atmosphere and neglected the impact of induction. Interestingly, while the bias has increased, there is still good correlation between the ensemble mean estimate and the ground truth during the first 100 s.

We take the six best precursors for each case, simulate turbines, and compare the precursor

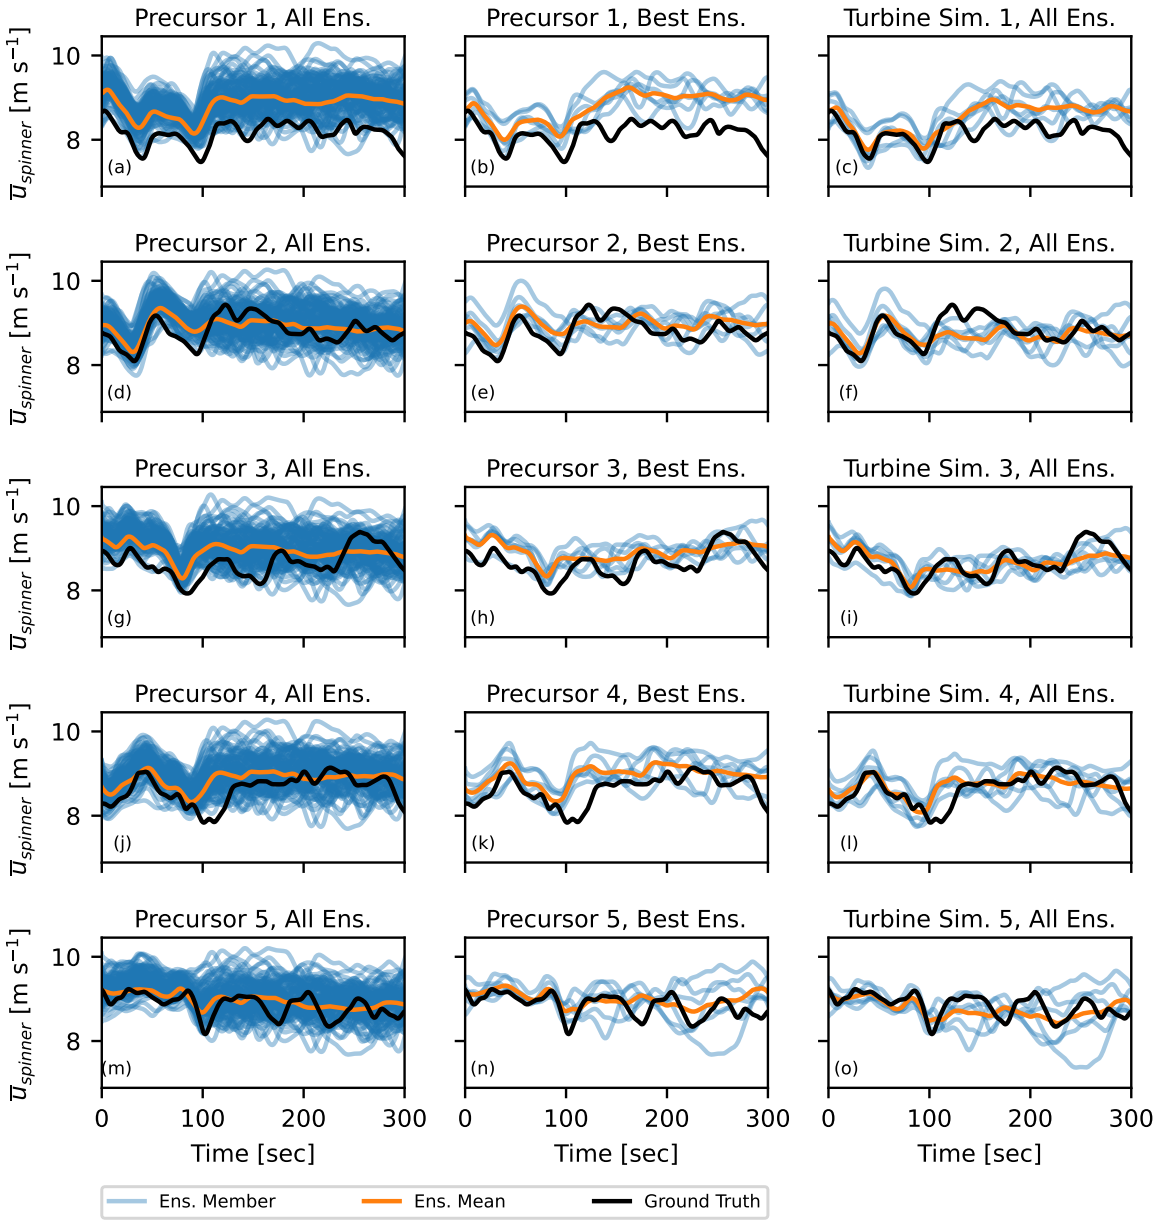

Figure 4. Spatially averaged spinner lidar measurements in the precursor simulations, with (left) all 100 ensemble members and (center) the six best ensemble members. (Right) Spatially averaged spinner lidar winds in the turbine-including simulations.

winds at the spinner lidar to the turbine-impacted winds (figure 4 right). While we selected the ensemble members for turbine simulations based off of turbine-free winds, we were unaware if the ensemble would be significantly impacted by the presence of a turbine. We find that the turbine-including ensemble responds well, lowering the six-member ensemble mean bias to 0.10 m s^{-1} across all cases in the first 100 s. The correlation of the six-member turbine-including ensemble over this period is 0.80.

Finally, we compare our estimated turbine dynamics (generator power and the Blade 1 root flapwise bending moment) to the truth (figure 5). We partition the turbine estimate window

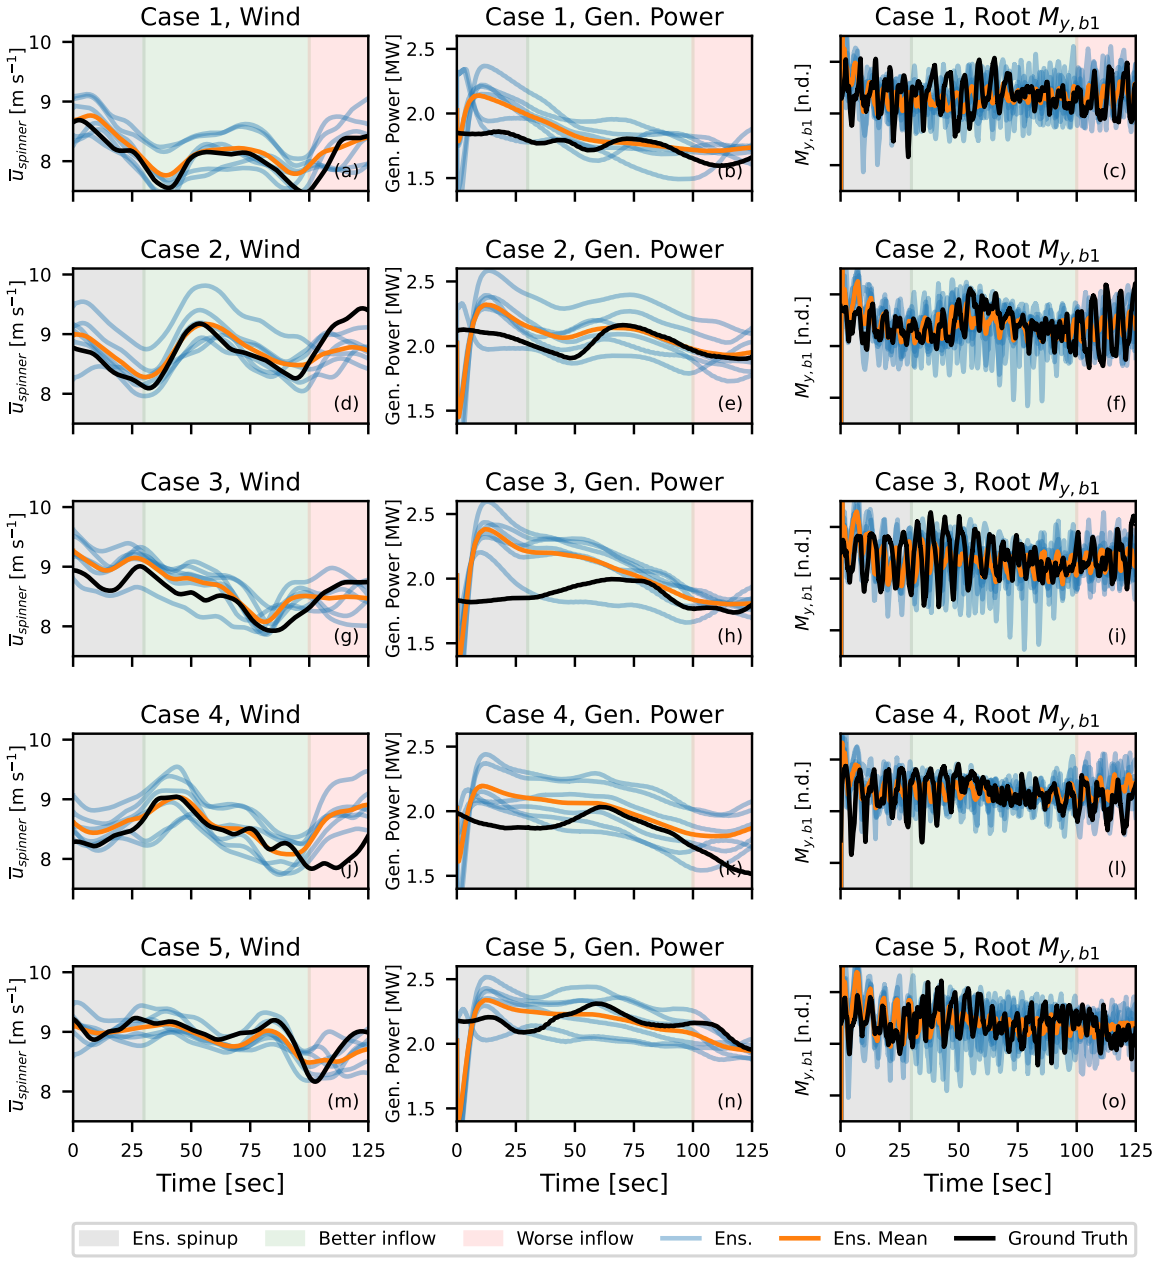

Figure 5. (Left) The six best spinner lidar wind reconstructions in the turbine-including simulations. The estimated and ground truth (center) generator power and (right) blade root flapwise bending moment $M_{y,b1}$.

into three phases. Turbine spin-up occurs in the estimate simulations during the first 30 s, and it is fruitless to compare estimated turbine behavior to true turbine behavior here. We denote 30–100 s as the period with better estimated inflow, during which we expect the best estimate of turbine dynamics. After 100 s, the inflow estimates decorrelate, so we omit turbine analysis here as well.

During the 30–100 s window, we estimate generator power and the blade root flapwise bending moment fairly accurately. For both quantities, the ground truth almost always falls within the spread of the six-member estimate. While the ensemble wind estimate is fairly accurate throughout the entire window, the generator power in all five cases only matches the ensemble estimate well in the latter half of the window. This suggests that perhaps the designation of turbine transients should be extended from 30 s to 60 s, though the decent correlation in the first half of the Case 2 power (figure 5e) muddles this conclusion. Unlike power, the ground truth bending moment agrees well with its ensemble mean estimate throughout the window, perhaps because this quantity responds more quickly to changing winds than power does. We note that the ensemble mean bending moment has a relatively small amplitude because of small phase offsets between individual estimates. All in all, this analysis of estimated turbine dynamics demonstrates that our flow reconstruction approach will estimate certain turbine behaviors better than others.

4. Discussion and conclusion

In this study, we applied a machine learning algorithm to read in synthetic wind observations and generate ensembles of initial conditions for large-eddy simulations. We tested the ability of this strategy to reconstruct time-varying inflow winds and turbine dynamics over a 10 min window. In case studies without turbines, we found that this approach can accurately reconstruct winds at approximately $1D$ upstream for the first 100 s after the measurement. Across all five case studies, ground truth winds always fell within the distribution of the 100 ensemble members during this period. We then applied the same machine learning algorithm to study simulations with turbines. The algorithm initializes the large-eddy simulations remarkably well, as estimated winds in turbine simulations at approximately $1D$ upstream showed a bias of 0.1 m s^{-1} and correlation of 0.80 during the first 100 s. When turning to turbine dynamics, the ensemble almost always spanned the bending moment of Blade 1 root flapwise bending moment and power, though the overall quality of the estimates varied between the two quantities. This study demonstrates that for short windows our flow reconstruction methodology may estimate certain turbine dynamics well while failing to capture other turbine characteristics accurately.

In order to improve the performance of the "guess-and-check" flow reconstruction strategy here, this work could be extended in a number of manners. For example, the LDM could be trained on turbine-containing atmospheres. This would help mitigate the bias in the LDM here due to induction effects, and it could potentially reduce the turbine spin-up transient, thereby prolonging the useful window of estimated turbine dynamics. Additional lines of inquiry could be pursued to address the bias present in the turbine-free estimates, to further reduce member-to-member spread and improve estimate confidence, or to improve how "best" precursors are selected for turbine simulations.

The results of this study will serve as a baseline for future dynamic inflow reconstruction efforts. Our strategy here relies on estimating accurate initial conditions given a snapshot of observations. These initial conditions could be used as initial guesses for data assimilation techniques [10] like 4DVAR or the ensemble Kalman filter, which further improve initial conditions by accounting for time histories of observations. Alternatively, our approach could also be coupled to a time-aware model-observation fusion technique like observation nudging, which in essence provides boundary conditions. While we are interested in flow reconstruction in the context of turbine model validation, we believe interest in this area will continue to grow,

as accurate inflow reconstruction is important for other aspects of the wind energy field such as flow control.

Acknowledgments

This work was authored in part by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) under Contract No. DE-AC36-08GO28308. Funding provided by the U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Wind Energy Technologies Office. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. Government purposes. The research was performed using computational resources sponsored by the Department of Energy's Office of Energy Efficiency and Renewable Energy and located at the National Renewable Energy Laboratory.

References

[1] Doubrawa P et al. 2020 Multimodel validation of single wakes in neutral and stratified atmospheric conditions *Wind Energy* **23** 2027–55
[2] Rybchuk A et al. 2023 Generating Initial Conditions for Ensemble Data Assimilation of Large Eddy Simulations with Latent Diffusion Models *Preprint* arXiv:2303.00836
[3] Rombach R, Blattmann A, Lorenz D, Esser P and Ommer B 2022 High-resolution image synthesis with latent diffusion models *Preprint* arXiv:2112.10752
[4] Zauner M, Mons V, Marquet O and Leclaire B 2022 Nudging-based data assimilation of the turbulent flow around a square cylinder *J Fluid Mech* **937** A38
[5] Evensen G 1994 Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics *J Geophys Res: Oceans* **99** 10143–62
[6] https://github.com/Exawind/amr-wind
[7] https://github.com/NREL/openfast-turbine-models/tree/master/IEA-scaled/NREL-2.8-127
[8] Churchfield M, Lee S and Moriarty P J 2012 A large-eddy simulation of wind-plant aerodynamics 50th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition 9–12 January 2012 Nashville, TN
[9] World Meteorological Organization 2021 Guidelines on ensemble prediction system postprocessing WMO-No. 1254 https://library.wmo.int/index.php?lvl=notice_displayid=21911.Y87FGOzMIUE
[10] Carrassi, A, Bocquet M, Bertino L and Evensen G 2018 Data assimilation in the geosciences: an overview of methods, issues, and perspectives *WIREs Climate Change* **9** e535