Response to Referee #1

> The manuscript presents a control methodology for achieving a prescribed wind-speed and wind-direction pair at a target height in large-eddy simulations of the atmospheric boundary layer by dynamically adjusting the geostrophic forcing vector. The manuscript is generally clear, and the proposed algorithm is presented in a way that is understandable. However, I have several major concerns about the framing, novelty, physical interpretation, and validation of the method.

We sincerely thank the referee for their thorough and insightful review. Their comments regarding the thermodynamic evolution of the boundary layer and the relevant literature on pressure-gradient controllers are highly accurate and will significantly strengthen the manuscript. Below, we address each concern point-by-point and outline our planned revisions.

Major comments

> 1. My main concern is not with the mathematical formulation of the controller itself, but with the suitability of using a geostrophic-forcing controller to obtain, or claim to obtain, a statistically steady convective boundary layer. In a convective boundary layer, the boundary-layer depth evolves in time due to surface buoyancy forcing and entrainment at the capping inversion. Adjusting the geostrophic forcing to match a target-height wind vector primarily controls the momentum balance, whereas the boundary-layer growth is thermodynamically driven and only indirectly coupled to the momentum controller.
>
>   This distinction is important because the manuscript frames DGN as a way of avoiding excessive transient boundary-layer growth by shortening the spin-up period. However, the results shown in the manuscript indicate that the inversion-height diagnostic continues to increase in time in both the nudged and non-nudged simulations. This is expected for a convective boundary layer with prescribed surface heat flux. The controller may reduce the time required for the target-height wind vector to converge, but it does not by itself produce a statistically steady convective boundary layer in the thermodynamic sense.
>
>   Therefore, there appears to be a mismatch between the controller formulation and the selected demonstration problem. The method is momentum-based, while the main complication emphasized for the convective case is the thermally driven growth of the boundary layer. The manuscript would benefit from a clearer separation between: (i) convergence of the target-height mean wind vector, and (ii) statistical stationarity of the boundary-layer thermodynamic and turbulence structure. Relevant work on statistically steady convective and stable boundary-layer simulations, such as Grylls et al. (2020) and Bon et al. (2024), should be discussed in this context.

Thank you for your comments and we completely agree. Our framing conflated the rapid achievement of a target momentum state with thermodynamic stationarity. In a convective boundary layer with a prescribed surface heat flux, the boundary layer will continuously deepen. DGN does not alter this thermodynamic reality. Rather, its primary benefit is that it accelerates the momentum convergence (spin-up), allowing the user to extract a usable, quasi-stationary window of data *before* the boundary layer grows excessively deep and interacts with the domain top.

Planned Revision: We will revise the Abstract and Introduction to clearly separate momentum convergence from thermodynamic stationarity. We will remove claims that DGN "arrests" boundary layer growth, clarifying instead that it allows the simulation to reach the target wind vector before the thermally driven growth necessitates a taller domain. We will also include the suggested references (Grylls et al., 2020; Bon et al., 2024) to contextualize true steady-state ABL methodologies.

> 2. A second major concern is that the manuscript does not sufficiently engage with prior work on feedback control of atmospheric-boundary-layer LES. The literature discussion mainly contrasts DGN with direct velocity nudging and trial-and-error selection of a constant geostrophic forcing. However, this misses important prior studies that already introduced controllers acting through the geostrophic forcing rather than by direct relaxation of the velocity field.
>
>   In particular, Sescu and Meneveau (2014) and Allaerts and Meyers (2015) should be discussed explicitly. Allaerts and Meyers (2015), for example, used a wind-angle controller to regulate the direction of the driving pressure gradient so that the hub-height wind direction remained aligned with the desired turbine-array geometry.
>
>   The present manuscript therefore should not frame the broad idea of controlling ABL LES through geostrophic or pressure-gradient forcing as entirely new. The more specific contribution appears to be a vector geostrophic-forcing controller that targets both windspeed magnitude and wind direction at a reference height.

We thank the referee for directing us to these critical references, which were a significant oversight in our literature review. We acknowledge that controlling the driving pressure gradient to achieve specific ABL conditions is not an entirely new concept.

Planned Revision: We will integrate Sescu and Meneveau (2014) and Allaerts and Meyers (2015) into the Introduction. We will reframe the novelty of DGN: rather than introducing the broad concept of geostrophic forcing control, DGN contributes a specific, closed-loop mechanism that explicitly incorporates the system's *current tendency* (Eq. 5) across both vector components to rapidly damp inertial oscillations.

> 3 .The manuscript controls not only the geostrophic-wind direction but also the geostrophicwind magnitude. While this provides an additional degree of freedom for matching a desired wind vector at the target height, it also changes the governing nondimensional state of the boundary layer.

> The structure of an ABL depends on nondimensional parameters such as the geostrophic drag coefficient (u_*/G), the Rossby number, stability parameters, and, where stratification or gravity waves are relevant, Froude-number-like parameters. Changing (G) during the simulation can therefore alter the nondimensional boundary-layer state, not merely accelerate convergence toward the same physical solution.
>
> This point is particularly important for the interpretation of Fig. 7. The comparison of turbulence structure between the nudged and non-nudged cases is not necessarily a comparison between two approaches to reaching the same boundary-layer state. It may instead be a comparison between two boundary layers evolving under different effective geostrophic forcings. The manuscript should either demonstrate that the final nondimensional states are equivalent or explicitly acknowledge that DGN changes the large-scale forcing and therefore the boundary-layer regime during the controlled transient.
>
> A stronger validation would compare DGN against a conventional simulation using the final converged geostrophic forcing inferred by the controller, or against an established wind-angle/geostrophic controller from the literature. This would separate the effect of the controller from the effect of solving a different forced ABL problem.

This is an excellent physical insight. By dynamically updating $\mathbf{G}$, the geostrophic drag coefficient (u_*/G) and the Rossby number evolve during the transient phase. Therefore, Figure 7 is indeed comparing a boundary layer under constant forcing to one experiencing time-varying forcing during its spin-up.

Planned Revision: We will explicitly acknowledge in Section 4.2 that the controller alters the non-dimensional boundary-layer regime during the transient phase. To provide a stronger validation, we will run an additional "A Priori" 3D benchmark simulation (namely, **Calibrated**, using a constant, calibrated forcing computed by eq. 16) using the *final, converged* geostrophic forcing inferred by the DGN controller. We will add a comparison to this benchmark to isolate the effect of the controller from the effect of solving a fundamentally different forced ABL problem.

> 4. The manuscript identifies the boundary-layer height using the altitude of maximum vertical gradient of potential temperature. This diagnostic is more accurately interpreted as a thermodynamic estimate of the inversion-layer height, rather than a general turbulence- or momentum-based boundary-layer depth.

We agree with this terminology correction.

Planned Revision: We will replace "Boundary Layer Height (BLH)" with "inversion-layer height" or "thermodynamic boundary layer depth" throughout the text and figures when referring to this specific diagnostic.

> 5. The 3D WRF-LES validation uses a vertical domain of only 1.2 km, with a Rayleigh damping layer applied over the upper 300 m. The manuscript itself reports that the

> diagnosed boundary-layer height begins to interact with the Rayleigh damping layer after approximately 18 h.
>
> This raises concerns about using the full 24 h simulation as validation of controller performance and turbulence preservation. If the boundary layer or inversion region interacts with the damping layer, the resulting statistics may be contaminated by the artificial upper-boundary treatment. The authors should either restrict the statistical analysis to the uncontaminated period, increase the vertical domain height, or provide additional evidence that the damping layer does not affect the reported conclusions.
>
> The concern is not simply that a 1.2 km domain is always inappropriate; rather, in this specific simulation, the authors' own results indicate interaction with the damping layer. This should be treated as a limitation of the validation case

The referee is correct to raise this concern. It was precisely because of this interaction that we restricted our statistical analysis of the turbulence structure (Figure 7c) to the 10–15 hour window, well before the inversion layer reaches the damping zone at 18 hours.

Planned Revision: We will make this methodological choice much more explicit in the text. We will clarify in Section 4.2 that all statistical validations of the turbulence structure are strictly confined to the uncontaminated 10–15 hour period to avoid artifacts from the Rayleigh damping layer.

Minor comments 1.

> 1. Lines 44–47: Is this sentence misplaced?

We will review the flow of this paragraph. The transition from the targeting problem to the historical mitigation strategies will be rewritten for better logical cohesion.

> 2. Figure 2: RMSE and standard deviation are used inconsistently. The manuscript defines the forcing convergence error (ε_G) as an RMSE of the controlled geostrophic wind relative to the analytical Ekman solution. However, the middle panel of Fig. 2 and its color bar refer to the standard deviation of the geostrophic wind. These are different quantities. RMSE relative to the analytical target measures convergence to the correct forcing, whereas standard deviation measures temporal variability about the mean. There is no mathematical justification for treating the two as equivalent unless the mean is exactly equal to the analytical target, which must be demonstrated rather than assumed. The authors should clarify whether panel b shows RMSE relative to the analytical forcing or the temporal standard deviation of the geostrophic forcing. The figure label, caption, and text should be made consistent

We apologize for this labeling error. Panel 2b displays the Root Mean Square Error (RMSE) relative to the analytical Ekman solution, *not* the temporal standard deviation. *Planned*

Revision: We will correct the y-axis label and color bar in Figure 2 to state "RMSE" for consistency with the text and equation.

> 3. The manuscript would be strengthened by comparing DGN with established wind-angle or geostrophic-forcing controllers rather than only with a fixed-(G) “NoNudge” case. The fixed-(G) baseline demonstrates that DGN is better than a naive forcing choice, but it does not establish superiority over the relevant state of the art.

We agree that comparing solely against a fixed-$\mathbf{G}$ baseline is limited. To demonstrate DGN's advantage over existing PI-based geostrophic controllers (e.g., Allaerts and Meyers, 2015), we will add a 1D column model comparison. This will clearly illustrate how DGN's explicit tendency term more effectively damps inertial oscillations than standard PI feedback. For the 3D case, we will add an "A Priori" benchmark (using the converged forcing derived through the method in Appendix A) to provide a stricter baseline.

> 4. Given the formulation of the controller, a neutral or conventionally neutral boundary layer may provide a cleaner validation case. Such cases are more directly tied to the momentum balance controlled by geostrophic forcing and avoid conflating target-height wind-vector convergence with thermally driven boundary-layer growth

We agree that a neutral boundary layer isolates the momentum balance and provides a purer theoretical test. We selected the convective case specifically to demonstrate DGN's practical operational benefit: achieving momentum convergence *before* unmanageable, thermally driven boundary layer growth occurs. We will clarify this rationale in Section 4.1 and note in the Conclusion that validating DGN across strictly neutral and stable regimes remains an important next step.

References:
[1] Grylls, T., Suter I., Reeuwijk, M.: Steady-State Large-Eddy Simulations of Convective and Stable Urban Boundary Layers, Boundary-Layer Meteorology, 175, 309–341, 2020
[2] Bon, T., Bayoán R., Meyers, J.: Stable Boundary Layers with Subsidence: Scaling and Similarity of the Steady State, Boundary-Layer Meteorology, 190:42, 2024
[3] Allaerts, D., Meyers, J.: Large eddy simulation of a large wind-turbine array in a conventionally neutral atmospheric boundary layer, Physics of Fluids, 27, 065108 (2015)
[4] Sescu, A., Meneveau, C.: A control algorithm for statistically stationary large-eddy simulations of thermally stratified boundary layers, Q. J. R. Meteorol. Soc. 140(683), 2017–2022, (2014).