




GNSS Zenith Wet Delay as a Boundary Layer Diagnostic: Regime-Dependent Turbulence Signatures From Large Eddy Simulation and Observations

 Gaël Kermarrec¹ , Tim Schrader¹, Xavier Calbet² , and Zhiguo Deng³ 
¹Institute for Meteorology and Climatology, Leibniz Universität Hannover, Hannover, Germany, ²AEMET, Madrid, Spain,

³Deutsches GeoForschungsZentrum GFZ, Potsdam, Germany

Key Points:

- large-eddy simulations (LES) show that GNSS-derived ZWD variance is synchronous with turbulent kinetic energy (TKE) and that the cutoff frequency is dynamically out of phase with boundary-layer-averaged wind speed during convection
- Three years of co-located Global Navigation Satellite Systems and Doppler LiDAR observations at Payerne reproduce the simulated coherence patterns under summer conditions, with longer lags between the spectral parameters and the boundary-layer variables than in the idealized LES
- A diurnal coherence framework based on Dynamic Time Warping quantifies the temporal coupling between spectral parameters and boundary-layer turbulence, avoiding the artifacts of correlation analysis on diurnally forced cycles

Correspondence to:

 G. Kermarrec,
kermarrec@meteo.uni-hannover.de

Citation:

Kermarrec, G., Schrader, T., Calbet, X., & Deng, Z. (2026). GNSS zenith wet delay as a boundary layer diagnostic: Regime-dependent turbulence signatures from large eddy simulation and observations. *Journal of Geophysical Research: Atmospheres*, 131, e2026JD047347. <https://doi.org/10.1029/2026JD047347>

Received 28 APR 2026

Accepted 12 JUN 2026

Author Contributions:

Conceptualization: Gaël Kermarrec
Data curation: Tim Schrader
Formal analysis: Gaël Kermarrec
Funding acquisition: Gaël Kermarrec
Investigation: Gaël Kermarrec, Tim Schrader
Methodology: Gaël Kermarrec
Project administration: Gaël Kermarrec
Software: Gaël Kermarrec, Zhiguo Deng

© 2026. The Author(s).

This is an open access article under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

Abstract Global Navigation Satellite Systems (GNSS) provide continuous measurements of zenith wet delay (ZWD), reflecting column-integrated atmospheric water vapor. While the slowly varying ZWD is routinely assimilated in numerical weather prediction, the rapid fluctuations on timescales of seconds to minutes that arise from boundary-layer turbulence remain underexploited. We test the hypothesis that the variance σ^2 and the cutoff frequency λ of the von Kármán model fitted to the ZWD power spectrum carry diagnostic information about boundary-layer state. We develop a diurnal coherence framework based on Dynamic Time Warping, extrema lags and phase-resolved slope ratios. The framework is calibrated on large-eddy simulations (LES) of an idealized convective boundary layer. In the simulation, σ^2 tracks turbulent kinetic energy (TKE) synchronously, λ leads TKE by half an hour and trails the integrated humidity and temperature variances by 1–1.5 hr, and the boundary-layer-averaged horizontal wind speed V_h is out of phase with both spectral parameters during convection. The framework is then applied without modification to 3 years of co-located GNSS and Doppler LiDAR observations at Payerne, Switzerland (249 clear-sky days, stratified by season). The summer subset reproduces the simulated patterns, with σ^2 tracking TKE and V_h – λ remaining out of phase. Lags are longer in the observations than in the LES, consistent with sources of variability absent from the idealized simulation. Bootstrap uncertainties on λ remain below 5% throughout the diurnal cycle. These results support the use of ZWD spectral parameters as a diagnostic of convective boundary-layer state, complementary to dedicated remote-sensing instruments.

Plain Language Summary Global Navigation Satellite Systems (GNSS) signals, used routinely for satellite positioning, are slightly delayed as they cross the moist atmosphere. Beyond the slow variations associated with weather forecasting, these delays exhibit rapid fluctuations on timescales of seconds to minutes that are caused by turbulent mixing of water vapor in the lowest kilometers of the atmosphere. We investigate whether two properties of these fluctuations, the total variance and the cutoff frequency of their power spectrum, carry quantitative information on the state of the atmospheric boundary layer. The difficulty is that all the variables of interest, including the wind speed, the humidity variance and the spectral parameters are dominated by a common diurnal cycle which inflates ordinary correlations and makes them physically meaningless. To address this, we develop a coherence analysis that compares the shape of diurnal cycles independently of amplitude and timing offsets. When tested on a high-resolution simulation of a convective boundary layer, the analysis recovers the expected dynamical relations between turbulence intensity, wind speed and the spectral parameters. Run on 3 years of co-located GNSS and wind LiDAR observations from Payerne, Switzerland, the same patterns emerge under summer conditions, while they disappear in winter when the diurnal cycle of turbulence is muted. Spectral parameters can be retrieved from existing GNSS networks with a relative uncertainty below 5% throughout the day, including at night. These results support the use of GNSS as a complementary observing system for boundary-layer dynamics in regions lacking dedicated meteorological instrumentation.

1. Introduction

As Global Navigation Satellite Systems (GNSS) microwave signals travel from satellites to ground-based receivers, they interact with the neutral atmosphere, experiencing propagation delays proportional to the integrated refractivity along the signal path (Bevis et al., 1992). The total tropospheric delay comprises two components: a hydrostatic delay depending primarily on surface pressure, and a wet delay sensitive to the distribution of water

Validation: Gaël Kermarrec
Visualization: Gaël Kermarrec
Writing – original draft: Gaël Kermarrec
Writing – review & editing:
 Tim Schrader, Xavier Calbet, Zhiguo Deng

vapor. The so-called zenith wet delay (ZWD) represents the vertically integrated contribution of water vapor to signal propagation delay. It typically ranges from a few centimeters in dry polar or continental conditions to over 30 cm in humid tropical environments (Crocetti et al., 2024). The relationship between ZWD and column-integrated water vapor is well established, enabling operational retrieval of precipitable water vapor (PWV) from GNSS networks (Duan et al., 1996; Qiu et al., 2020). These PWV estimates are now routinely assimilated into numerical weather prediction (NWP) systems, improving forecasts of precipitation and severe weather events (Guerova et al., 2016).

While the slowly varying ZWD (timescales of hours to days) has been extensively studied, the rapid fluctuations superimposed on this mean signal remain largely unexploited for atmospheric science. These fluctuations, with amplitudes of a few millimeters and timescales from seconds to minutes, originate primarily from turbulent variations in the wet refractivity field within the atmospheric boundary layer (Kermarrec, Czerwonka-Schröder, & Holst, 2023; Kermarrec, Calbet, et al., 2025). The wet refractivity depends jointly on water vapor partial pressure and temperature (Bevis et al., 1992), so that turbulent eddies generate ZWD fluctuations through moisture and thermodynamic pathways simultaneously as they transport heat and water vapor across the column. Other atmospheric processes also contribute to short-term ZWD variability, in particular convective events and short-lived intense precipitation that can produce non-stationary refractivity excursions (Ma et al., 2022; Nahmani et al., 2019), and high-frequency gravity waves whose vertical-wavenumber spectra have been extensively characterized in the free troposphere and middle atmosphere (Fritts & Alexander, 2003). The present analysis is restricted to the turbulent component by combining clear-sky filtering with the inherent selectivity of the von Kármán spectral fit, which rejects non-stationary or non-Kolmogorov contributions through the Whittle likelihood residual test (Sykulski et al., 2019).

The power spectral density of these fluctuations follows a low-frequency plateau transitioning to Kolmogorov scaling ($f^{-8/3}$ for vertically integrated quantities, where f is the temporal frequency) at higher frequencies (Lilly et al., 2017; Tatarski et al., 1961; Wheelon, 2001). Two parameters characterize this spectrum: the variance σ^2 , quantifying fluctuation intensity, and the cutoff frequency λ , marking the transition between the energy-input region and the inertial subrange. These parameters have been estimated from GNSS observations in various climate zones (Kermarrec & Schön, 2020; Kermarrec, Calbet, et al., 2025), but their physical relationship to boundary-layer dynamics remains poorly understood.

The motivation for further characterizing ZWD spectral parameters is distinct from the established assimilation of column-integrated PWV in NWP (Guerova et al., 2016; Jones et al., 2020). Operational assimilation exploits the slowly varying total water content and does not require knowledge of boundary-layer turbulence. A separate question is whether the rapid spectral content of ZWD can be used as a real-time diagnostic of boundary-layer state, in a manner analogous to Doppler LiDAR or scintillometer measurements but exploiting an existing global infrastructure of more than 30,000 GNSS stations (International GNSS Service, 2023). Unlike Doppler LiDAR, radiosondes or flux towers (Foken, 2008; Pearson et al., 2009; Seidel et al., 2010), GNSS receivers are already deployed, continuously operating, and require no specialized meteorological expertise to maintain. This complementarity, rather than redundancy, with respect to current NWP applications motivates the present study. We address three questions. First, which boundary-layer variables control the ZWD variance σ^2 and cutoff frequency λ in a controlled simulated environment? Second, does the temporal coupling between these spectral parameters and boundary-layer variables persist in real atmospheric conditions, and how does it depend on the dynamical regime? Third, can the same methodology be transferred from idealized simulations to multi-year observations without methodological adjustment?

To answer these questions, we generate synthetic ZWD time series from LES of an idealized diurnal cycle (Zus et al., 2025), and we develop a diurnal coherence approach based on Dynamic Time Warping (DTW) to quantify the temporal coupling between spectral parameters and boundary-layer variables. This methodology is needed because correlation analysis applied to variables that share a common diurnal forcing produces inflated coefficients that do not reflect physical coupling (van Nes et al., 2015). Observations from Payerne, Switzerland, where GNSS and Doppler wind LiDAR are co-located, provide a 3-year validation data set stratified by season.

The paper is organized as follows. Section 2 describes the LES configuration, the synthetic ZWD generation, the spectral parameter estimation, and the diurnal coherence framework. Section 3 presents the LES coherence results. Section 4 transfers the framework to the multi-year Payerne observations through case studies and

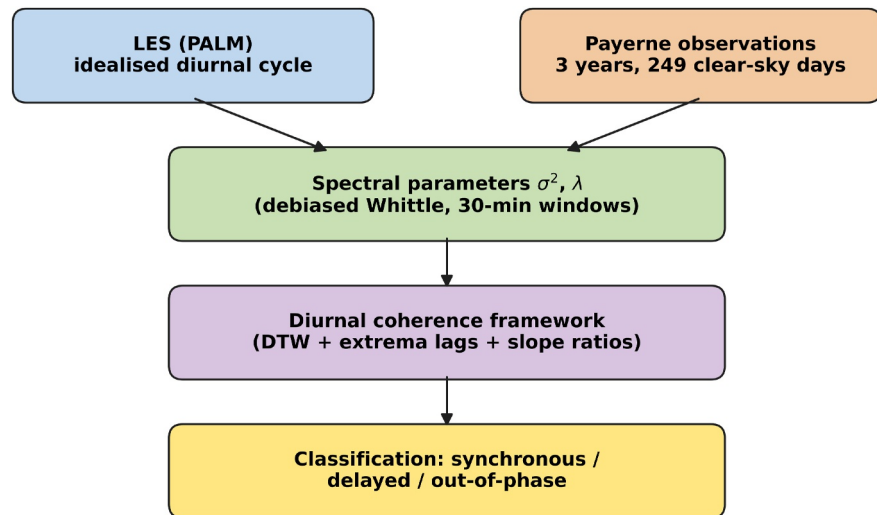


Figure 1. Workflow of the analysis. The large-eddy simulations (LES) pipeline produces synthetic ZWD time series at 50 randomly distributed locations. Spectral parameters σ^2 and λ are estimated by debiased Whittle maximum likelihood on 30-minute sliding windows. Boundary-layer predictors are computed in parallel from the LES output. The diurnal coherence framework is then applied to quantify the temporal coupling between predictors and spectral parameters. The same framework is transferred to the multi-year Payerne observations.

aggregated statistics with seasonal stratification. Section 5 discusses the physical interpretation, the implications for boundary-layer monitoring, and the limitations. Section 6 summarizes the main findings and outlines future directions.

2. Methodology

The methodology aims to characterize the temporal coupling between ZWD spectral parameters and boundary-layer variables, in a way that is robust to the shared diurnal forcing affecting all variables under consideration. It proceeds in four steps. First, we generate synthetic ZWD time series from LES with fully characterized boundary-layer dynamics. As a turbulence-resolving numerical approach, LES explicitly resolves the energy-containing eddies of the boundary layer while parameterizing the smaller subgrid-scale motions, and it has become a standard tool for boundary-layer process studies (Maronga et al., 2020). Second, we extract spectral parameters from the turbulent fluctuations using debiased Whittle maximum likelihood estimation. Third, we apply a diurnal coherence framework based on DTW, extrema lags and phase-resolved slope ratios to quantify the temporal coupling between spectral parameters and boundary-layer variables. Fourth, we transfer this framework to multi-year observations from Payerne, Switzerland, with seasonal stratification. Each step is described below, and the overall workflow is summarized in Figure 1.

2.1. LES Setup

We use the Parallelized Large-Eddy Simulation Model (PALM) version 6.0 to simulate boundary layer evolution during a late-summer diurnal cycle (Maronga et al., 2020). As a turbulence-resolving model, PALM solves the filtered Navier-Stokes equations in the Boussinesq approximation with prognostic equations for potential temperature, specific humidity, and subgrid-scale TKE. The simulation employs two-way nesting with a parent domain of $4 \times 4 \text{ km}^2$ at 10 m resolution. The vertical extent reaches 1,200 m with grid stretching above 450 m. Lateral boundaries are cyclic in the zonal direction, with Rayleigh damping above 1,650 m preventing spurious wave reflections.

Initial conditions represent late-summer mid-latitudes: surface pressure 1,000 hPa, surface potential temperature 295 K, and geostrophic wind 4 m s^{-1} from the west. A stable stratification (0.7 K per 100 m above 800 m) caps the boundary layer. Following a 24-hour spin-up, the simulation spans 15 hr (09:00–24:00 local time), capturing the full transition from morning convective development through evening stabilization. Radiative forcing is computed using the Rapid Radiative Transfer Model for GCMs (Clough et al., 2005; Maronga et al., 2020) under

clear-sky conditions. Based on the surface heat flux evolution, we define 17:00 local time as the day/night transition, corresponding to the onset of stable stratification as convective forcing diminishes.

Three-dimensional fields of velocity components (u , v , w), potential temperature (θ), specific humidity (q), and TKE are output at 30-minute intervals. High-frequency output (1 Hz) along vertical columns enables synthetic ZWD generation at temporal resolution matching GNSS observations.

LES of the convective boundary layer at 10 m horizontal resolution over a $4 \times 4 \text{ km}^2$ domain are computationally demanding, with each simulation requiring a substantial allocation of supercomputing resources. To extract robust statistical properties of the synthetic ZWD without running multiple simulations, we sample 50 randomly distributed vertical columns within the LES domain, each providing an independent ZWD time series at the 1 Hz output cadence. This intra-domain ensemble sampling strategy follows the approach validated by Ralves et al. (2022) for virtual atmospheric measurements within LES domains, and substitutes ensemble sampling for repeated runs. The 50 columns are sufficient to converge the spectral statistics of σ^2 and λ to within their estimation uncertainty, as verified by sub-sample bootstrap tests.

2.2. Synthetic ZWD Generation

The wet refractivity N_w depends on both water vapor partial pressure e and temperature T through the Smith and Weintraub formulation (Bevis et al., 1992):

$$N_w = k_2' \frac{e}{T} + k_3 \frac{e}{T^2}, \quad (1)$$

where $k_2' = 22.1 \text{ K Pa}^{-1}$ and $k_3 = 3.739 \times 10^5 \text{ K}^2 \text{ Pa}^{-1}$. The two terms reflect distinct physical pathways: the first captures the polar character of the water molecule and is approximately linear in e , whereas the second represents the dipole orientation effect that becomes dominant at lower temperatures. The ZWD is the path-integrated wet refractivity along the vertical (Zus et al., 2025),

$$\text{ZWD} = 10^{-6} \int_0^{z_{\text{top}}} N_w(z) dz. \quad (2)$$

Both temperature and humidity fluctuations therefore contribute to ZWD variability through Equation 1, and a turbulence-driven analysis of σ^2 and λ inherently mixes the moisture and thermodynamic pathways. The relative weight of the two contributions, and the role of other refractivity-relevant quantities such as gravity-wave-induced fluctuations or transient convective excursions, is discussed in Section 5.2. We extract ZWD time series from 50 randomly distributed locations within the LES domain to ensure statistical robustness while keeping the computational cost tractable; this approach has been shown to recover the same spectral statistics as a fully resolved spatial sampling (Kermarrec, Calbet, et al., 2025). Because the LES domain extends only to the top of the simulated boundary layer (1,200 m), the synthetic ZWD computed here corresponds to the boundary-layer contribution rather than to the full tropospheric column. This is a deliberate restriction: the high-frequency spectral content targeted by our analysis is dominated by boundary-layer turbulence, but residual contributions from the free troposphere are not captured by the simulation.

2.3. Spectral Parameter Estimation

Following our previous works, we assume that the ZWD power spectrum follows a von Kármán model (Kermarrec, Calbet, et al., 2025; Wheelon, 2001):

$$W(f) = \frac{\sigma^2 \lambda^{5/3}}{c_\lambda (f^2 + \lambda^2)^{4/3}} \quad (3)$$

where f is the frequency (Hz), σ^2 is the variance [mm^2], λ is the cutoff frequency (Hz), and c_λ is a dimensionless normalization constant. This spectral form is characteristic of temporal observations: while spatial spectra drop at low wavenumbers beyond the outer scale, temporal spectra saturate to a plateau at low frequencies to accommodate finite variance (Stull, 1988).




Figure 2. Idealized von Kármán spectrum of ZWD fluctuations showing the two spectral parameters: the variance σ^2 (area under the curve) and the cutoff frequency λ (spectral break). At frequencies below λ , the spectrum exhibits a flat plateau characteristic of the energy-input range for time series. Above λ , the spectrum follows a Kolmogorov $f^{-8/3}$ decay in the inertial subrange, steeper than the classical $-5/3$ slope due to vertical path integration.

The variance σ^2 quantifies the total power of ZWD fluctuations and reflects the integrated magnitude of humidity variability along the signal path. Physically, σ^2 responds to turbulent mixing intensity: stronger convection generates larger humidity fluctuations, increasing ZWD variance.

The cutoff frequency λ delimits two spectral regimes, as illustrated in Figure 2. At frequencies below λ , the spectrum is flat, corresponding to the energy-input range where large-scale atmospheric instabilities inject variance into the turbulent cascade (Stull, 1988). At frequencies above λ , the spectrum follows Kolmogorov scaling ($f^{-8/3}$ for vertically integrated quantities), corresponding to the inertial subrange where energy cascades toward smaller scales without external forcing. The cutoff frequency thus characterizes the temporal scale at which energy injection transitions to cascade dynamics. Under Taylor's frozen-flow hypothesis, λ relates to advection velocity and the outer scale of turbulence: faster advection “compresses” the temporal signature of turbulent structures, shifting λ toward higher frequencies. Variations from this spectrum can be found in, for example, Lilly et al. (2017).

Parameters are estimated using a debiased Whittle maximum likelihood estimator (Sykulski et al., 2019), processing 30-minute windows with 10-second sliding steps. The debiased Whittle approach removes the small-sample bias of the standard Whittle pseudolikelihood while retaining its

$O(n \log n)$ computational cost. Outlier batches whose periodogram residuals exceed three standard deviations from the fitted von Kármán form are excluded, affecting less than 1% of data (Kermarrec, Calbet, et al., 2025). This residual test is conservative: spectra dominated by gravity-wave resonances or by non-stationary convective transients deviate systematically from the von Kármán plateau-cascade form and are filtered out at the fit stage, in addition to the day-level clear-sky filtering described in Section 2.6.

Estimation uncertainty is quantified by bootstrap resampling. For each 30-minute window, we resample with replacement the 10-second batches contributing to the periodogram, refit the von Kármán model, and repeat the operation 200 times. The resulting standard error σ_λ is propagated to daily medians and to the diurnal cycles of λ shown in Section 4, where the magnitude of the relative uncertainty is reported and discussed.

2.4. Predictor Variables From LES

We select four predictor variables based on physical relevance and statistical independence:

Integrated humidity variance, $\int q^2$: the variance of vertically integrated humidity fluctuations, where $q'(z, t) = q(z, t) - \langle q(z, t) \rangle_{\text{horizontal}}$. This quantity directly parallels the σ^2 from ZWD, which integrates wet refractivity.

Integrated temperature variance, $\int \theta^2$: analogous to humidity but for potential temperature fluctuations, capturing buoyancy-driven turbulence.

Mean TKE, $\langle \text{TKE} \rangle$: boundary-layer-averaged TKE, providing a bulk measure of turbulence intensity.

Boundary-layer-averaged horizontal wind speed, $V_h = \sqrt{\langle u \rangle^2 + \langle v \rangle^2}$: the magnitude of mean horizontal wind, that is, averaged over the boundary layer height. Under Taylor's frozen-flow hypothesis, advection velocity controls how temporal fluctuations relate to spatial structure.

In the following, we may loosely refer to boundary-layer-averaged quantities by their variable name, omitting the averaging notation. Although these four predictors are not statistically independent, given that they all reflect the convective state of the boundary layer, the diurnal coherence framework introduced in Section 2.5 does not rely on independence. It treats each predictor-target pair separately and evaluates the temporal coupling without combining predictors in a regression model. We deliberately exclude the individual wind components $\langle u \rangle$ and $\langle v \rangle$ as separate predictors because the simulation is driven by a westerly geostrophic forcing, leading to $r(\langle u \rangle, V_h) = 1.00$ throughout the run, which would introduce an exact collinearity if the components were retained. The remaining predictors retain pairwise Pearson correlations between 0.74 and 0.99, but these are not

used as a basis for inference: they are noted only to motivate the avoidance of multivariate regression and the choice of a pairwise coherence framework instead. All variables are vertically averaged or integrated from the surface to the boundary layer height z_i , which is an output of the LES.

2.5. Diurnal Coherence Framework

The diurnal evolution of all variables under consideration is dominated by a shared forcing: the surface energy flux drives turbulent mixing, modulates wind shear and homogenizes humidity profiles in synchrony. On three to six points sampled within a single phase of the convective cycle, the Pearson and Spearman correlation coefficients reach values close to unity by virtue of this shared monotonicity, regardless of any underlying physical coupling. Inflated correlation coefficients in the presence of common forcing are a recognized difficulty in climate causal analysis (van Nes et al., 2015). The opposite failure mode is observed on the multi-year Payerne data set: the median intra-day Pearson correlations are $r(\lambda, V_h) = -0.06$ and $r(\sigma^2, \text{TKE}) = +0.14$ across the 249 clear-sky days, in contrast with $r = -0.99$ and $r = +0.99$ obtained for the same pairs from the LES. The two values bracket the same physical pair but differ by an order of magnitude, indicating that linear correlation is unreliable as a diagnostic of the underlying relationship. The mathematical reasons why are summarized in Appendix Appendix A; in short, the Pearson coefficient is sensitive to amplitude differences and is reduced by physical lags, so that it confounds the strength of the coupling with the timing of the response.

We adopt instead a coherence framework based on three complementary metrics. The DTW distance is the principal metric: it compares the shape of two diurnal cycles after standardization, allowing for a non-linear monotonic alignment of the time axis so that physical lags of a few hours are accommodated rather than penalized. The DTW distance is small when the cycles share a common shape, and the alignment itself returns the optimal lag as a separate output. Two local measures complete the framework: the lag between principal extrema and the phase-resolved slope ratio, which inform respectively on the timing of the diurnal markers and on the intra-phase covariation of the two series. The three metrics are described in turn below. Their consistency provides an internal robustness check that no single metric can supply on its own. The DTW framework has been used as a similarity metric for time series in geosciences for two decades (Berndt & Clifford, 1994), and was recently applied with Monte Carlo significance testing to the classification of GNSS displacement series in tectonic settings (Kumar et al., 2022), which is the closest published precedent for the present use.

2.5.1. Dynamic Time Warping Distance

For two time series $x(t)$ and $y(t)$ of equal length N , normalized to zero mean and unit variance, the DTW distance is defined as

$$d_{\text{DTW}}(x, y) = \frac{1}{N} \min_{\pi} \sum_{(i,j) \in \pi} |x_i - y_j|, \quad (4)$$

where π is a monotonic alignment between the indices of the two series, constrained to match the endpoints (Berndt & Clifford, 1994). The minimization is solved by dynamic programming in $O(N^2)$ operations. The resulting distance is sensitive to shape similarity independently of residual amplitude differences and modest temporal offsets. The mean optimal lag between the two series is defined as $\Delta t_{\text{DTW}} = \overline{(i - j)_{\pi^*}} \cdot \Delta t$, where π^* is the optimal alignment and Δt the sampling interval. With this definition, a *negative* Δt_{DTW} indicates that the first series leads the second, since alignment indices of the leading series are mapped to later indices of the trailing series.

We classify each pair of variables on the basis of d_{DTW} and Δt_{DTW} into three regimes. A pair is *synchronous* when $d_{\text{DTW}} < 0.30$ and $|\Delta t_{\text{DTW}}| \leq 0.5$ h, indicating a strong shape similarity without significant lag. It falls in the *synchronous with delay* regime when $d_{\text{DTW}} < 0.60$ and $|\Delta t_{\text{DTW}}| \leq 2$ h, marking a preserved similarity with a moderate physical delay. Beyond $d_{\text{DTW}} \geq 0.60$, the pair is classified as *out of phase*, reflecting a shape dissimilarity incompatible with direct coupling. The thresholds were calibrated on the idealized LES and applied without modification to the Payerne observations, ensuring the absence of post-hoc adjustment.

The robustness of the qualitative classification to the precise choice of thresholds was verified by independently varying the synchronous boundary in $\{0.25, 0.30, 0.35\}$ and the synchronous-with-delay boundary in

{0.55, 0.60, 0.65}, that is over an interval of $\pm 17\%$ around the values used here. Across the resulting 3×3 grid, the classifications of the eight LES pairs and the four summer Payerne pairs change for at most one pair (the borderline case $\langle \text{TKE} \rangle - \lambda$ in summer Payerne, with $d_{\text{DTW}} = 0.61$, oscillates between synchronous-with-delay and out-of-phase as the upper threshold crosses 0.60). The remaining classifications are stable. The conclusions of the analysis therefore do not depend on a particular choice of threshold within physically reasonable bounds.

2.5.2. Extrema Lags and Phase-Resolved Slope Ratios

The DTW distance summarizes the global similarity between two diurnal cycles. It is complemented by two local measures that inform on the nature of the coupling. The peak lag Δt_{peak} is defined as $t_{\text{peak},x} - t_{\text{peak},y}$, the difference between the time of the maximum of the predictor x and the time of the maximum of the target y , with both extrema constrained to the 06–22 hr window to avoid boundary effects associated with the day-night cycle. An analogous trough lag Δt_{trough} is defined on the minima. The lags are wrapped to the interval $[-12, +12]$ h to remove ambiguities for extrema near midnight. With this definition, a negative value indicates that the predictor leads the target. The same sign convention applies as for Δt_{DTW} .

The phase-resolved slope ratio quantifies the local covariation of the two series within diurnal sub-intervals defined a priori from the convective phenomenology. For each phase p with bounds $[t_{p,\text{start}}, t_{p,\text{end}}]$, we compute the normalized slope of the predictor as

$$s_x^{(p)} = \frac{x(t_{p,\text{end}}) - x(t_{p,\text{start}})}{(t_{p,\text{end}} - t_{p,\text{start}})A_x}, \quad (5)$$

where $A_x = \max x - \min x$ is the diurnal amplitude of the predictor. An analogous slope $s_y^{(p)}$ is defined for the target. The ratio $r^{(p)} = s_x^{(p)}/s_y^{(p)}$ is positive when the two variables evolve in the same direction during the phase and negative when they evolve in opposite directions; its modulus quantifies the relative rate of variation. For the LES analysis, the phases retained are Growth (9–12 hr), Peak (12–15 hr), Decay (15–19 hr) and Plateau (19–21.5 hr), reflecting the convective phenomenology of the simulated cycle. For the Payerne observations, which cover the full 24-hour cycle, an additional Pre-dawn phase (0–6 hr) is added.

2.5.3. Significance Testing of the Extrema Opposition Ratio

The extrema-based analysis of the Payerne observations quantifies the coupling between λ and V_h by the fraction of extrema pairs in opposition, that is the pairings between a λ maximum and a V_h minimum, or between a λ minimum and a V_h maximum, within a given temporal tolerance τ . A fraction larger than 0.5 indicates a preferentially anti-phased behavior. We assess the statistical significance of this fraction by a Monte Carlo permutation test. For each clear-sky day, the times of the extrema detected in V_h are replaced by independent uniform draws within the 06–22 hr window; the number and the type of extrema are preserved, while the extrema of λ are kept fixed. This procedure preserves the detection structure while destroying any temporal coupling between the two series. The opposition ratio is recomputed for each permutation and the operation is repeated 5,000 times. The resulting null distribution provides a one-sided p -value as the fraction of permutations whose ratio exceeds the observed value. A standardised score $Z = (\hat{r} - \mu_0)/\sigma_0$, where \hat{r} is the observed ratio and μ_0, σ_0 are the mean and standard deviation of the null distribution, supplies a comparable measure of departure from the null. This approach is conceptually analogous to that used by Kumar et al. (2022) to assess the significance of DTW-based GNSS classifications.

2.5.4. Sensitivity to Detection Parameters

The extrema-based analysis depends on two parameters: a prominence threshold p expressed as a fraction of the daily amplitude, and a temporal tolerance τ for the pairing of extrema between the two series. To rule out the dependence of the result on a particular choice of these parameters, we systematically evaluate the opposition ratio on the grid $p \in \{0.15, 0.20, 0.25, 0.30, 0.35\}$ and $\tau \in \{0.5, 1.0, 1.5, 2.0, 2.5, 3.0\}$ h. The stability of the ratio across the grid serves as a complementary robustness criterion to the permutation test.

2.6. Observational Data

We use observations from Payerne, Switzerland (46.82°N, 6.95°E, 491 m elevation) which hosts the MeteoSwiss aerological station with co-located GNSS receiver and Doppler wind LiDAR (Brocard et al., 2013; Martucci et al., 2010). The Swiss Plateau location experiences well-defined diurnal boundary layer cycles during fair-weather conditions, making it ideal for validating LES-derived relationships.

ZWD is estimated from dual-frequency GNSS observations using precise point positioning with ambiguity resolution (Kermarrec, Calbet, et al., 2025). The spectral parameters λ and σ^2 are extracted using the same methodology applied to LES data, ensuring consistency.

2.6.1. Doppler Wind Lidar

A Halo Photonics StreamLine XR scanning Doppler LiDAR provides vertical profiles of radial velocity at 30-minute resolution, spanning 100–3,000 m altitude with 50 m vertical resolution. The LiDAR operates in Doppler Beam Swinging mode, yielding the three wind components (u, v, w). TKE is computed from the velocity variances within each 30-min window:

$$\text{TKE}(z) = \frac{1}{2}(\sigma_u^2(z) + \sigma_v^2(z) + \sigma_w^2(z)). \quad (6)$$

Data gaps due to low aerosol backscatter are filled using linear interpolation across altitude levels before computing layer averages. Days with fewer than 40 valid time steps (out of 48) are excluded; outliers are removed by requiring $V_h < 10 \text{ m s}^{-1}$ to discard precipitation and low-SNR retrievals. For comparison with GNSS, boundary-layer-averaged quantities $\langle \text{TKE} \rangle$ and $V_h = \sqrt{\bar{u}^2 + \bar{v}^2}$ are computed by simple vertical averaging over the 100–850 m range. This upper limit reflects the increasing proportion of missing data and outliers above 850 m at Payerne under clear-sky conditions, and still captures the well-mixed convective layer while limiting free-tropospheric contamination.

2.6.2. Cloud Cover

Cloud cover information from the Payerne meteorological station provides context for interpreting humidity-related variability (Wacker et al., 2015). A Vaisala CL51 ceilometer (910 nm, 10 m vertical and 16 s temporal resolution) measures cloud base height and fractional cloud cover up to 15 km. Total cloud cover (0–8 oktas) is resampled to 30 min to match the spectral analysis cadence; data are excluded when co-located rain gauges detect precipitation or when backscatter profiles indicate full attenuation. The clear-sky filter ($\overline{\text{CC}} \leq 2$ oktas) used in Section 4.1 relies on this product.

3. LES Results

We first examine the diurnal evolution of the LES variables and the GNSS-derived spectral parameters, and we then apply the diurnal coherence framework introduced in Section 2.5 to quantify the temporal coupling between predictors and targets.

3.1. Diurnal Evolution of LES Variables

Figure 3 presents the diurnal evolution of the four boundary-layer predictors and the two GNSS-derived spectral parameters obtained from the LES. All quantities shown in this figure are derived from the simulation, and their interpretation as observable analogs is deferred to Section 4. The integrated humidity and temperature variances $\int \overline{q^2}$ and $\int \overline{\theta^2}$ follow pronounced convective cycles, peaking around 14–16 hr when surface heating maximizes buoyancy-driven mixing, and decaying rapidly after sunset. The boundary-layer-averaged TKE tracks this pattern closely. The boundary-layer-averaged horizontal wind speed V_h exhibits the opposite evolution during the convective phase, with minimum values around 3.6 m s^{-1} near the convective maximum when turbulent mixing reduces near-surface shear, followed by a gradual increase to 4.5 m s^{-1} in the evening as the boundary layer stabilizes and decouples from surface friction. This late-evening recovery is a feature of the idealized PALM configuration, in which a constant westerly geostrophic forcing of 4 m s^{-1} takes over once the surface flux declines. We return to its consequences for the interpretation of the V_h -related coherence in Section 3.2.

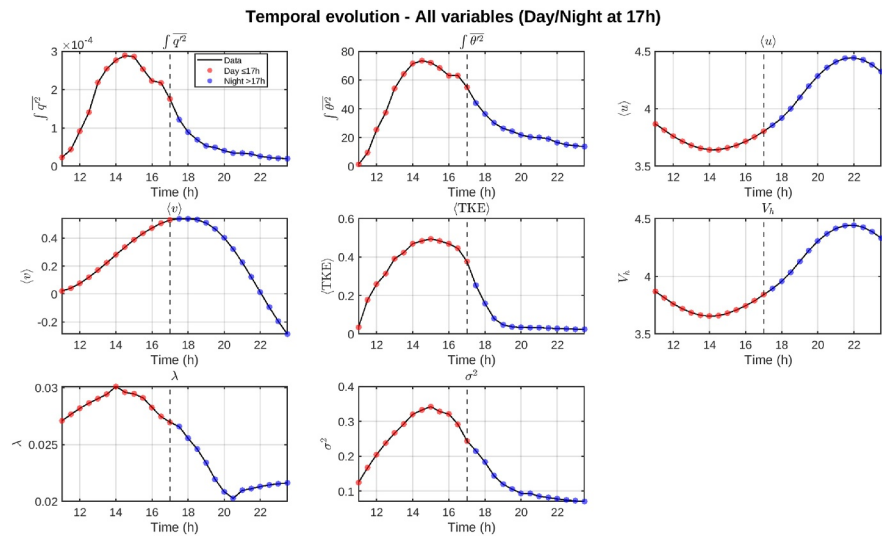


Figure 3. Diurnal evolution of the four boundary-layer predictors ($f \overline{q^2}$, $f \overline{\theta^2}$, $\langle \text{TKE} \rangle$, V_h) and the two GNSS-derived spectral parameters (λ , σ^2) computed from the large-eddy simulations. Red markers: daytime (≤ 17 h LT). Blue markers: nighttime (> 17 h LT). The vertical dashed line marks the day-night transition. The late-evening rebound of V_h reflects the idealized geostrophic forcing of the simulation and is not representative of observed clear-sky conditions.

The GNSS-derived parameters respond distinctly. The variance σ^2 tracks the bulk turbulence intensity, with a diurnal amplitude that mirrors that of $\langle \text{TKE} \rangle$. The cutoff frequency λ shows elevated values during daytime convection, reflecting the rapid turnover of buoyancy-driven structures, and decreases toward lower nocturnal values as fluctuations become more persistent under stable stratification. The bootstrap standard error of λ remains below 5% of its central value throughout the cycle, with values around 2% during the convective phase and 3%–4% in the early morning when λ approaches its lowest values, indicating that the diurnal variability of λ is well above the estimation uncertainty.

3.2. Pairwise Diurnal Coherence

We apply the diurnal coherence framework introduced in Section 2.5 to the eight pairs formed by the four predictors $f \overline{q^2}$, $f \overline{\theta^2}$, $\langle \text{TKE} \rangle$ and V_h with the two spectral targets λ and σ^2 . Figure 4 displays the diurnal cycles of each pair after standardization to zero mean and unit variance, with the four phases color-shaded in the background and the principal extrema marked. Each panel reports the DTW distance, the optimal DTW lag and the resulting classification.

The eight pairs split cleanly into three regimes. Six involving turbulence-related predictors fall in the synchronous or synchronous-with-delay categories; the remaining two, both involving V_h , classify as out of phase.

The strongest coherence appears for $\langle \text{TKE} \rangle$ – σ^2 ($d_{\text{DTW}} = 0.14$, near-zero lag), reflecting the direct physical chain in which turbulent mixing drives humidity and temperature fluctuations whose path-integrated variance is σ^2 . Consistent with this, $f \overline{q^2}$ – σ^2 and $f \overline{\theta^2}$ – σ^2 are also synchronous (distances 0.20 and 0.21, lags within 0.5 hr), since the path integrals of the two refractivity-relevant variances are modulated by the same mixing intensity. The pair $\langle \text{TKE} \rangle$ – λ joins the synchronous regime ($d_{\text{DTW}} = 0.24$, lag +0.5 h), with λ leading $\langle \text{TKE} \rangle$ by half an hour. The spectral break responds to the same convective forcing as bulk turbulence intensity, with a small lead because λ encodes the energy-input/inertial-subrange transition that emerges as soon as the first organized structures appear, slightly before the bulk intensity has fully matured. The scalar variances follow with a modest delay— $f \overline{q^2}$ – λ at $d_{\text{DTW}} = 0.33$, lag +1.0 h and $f \overline{\theta^2}$ – λ at 0.38, +1.5 h—consistent with the same reading: the spectral transition emerges with the first coherent eddies, while the path-integrated scalar variances accumulate over a longer timescale.

By contrast, both pairs involving V_h classify as out of phase (V_h – λ : $d_{\text{DTW}} = 1.46$, lag –4.5 h; V_h – σ^2 : 1.17, +4.5 h). Interpretation requires care, because the simulated V_h cycle increases through the second half of the day

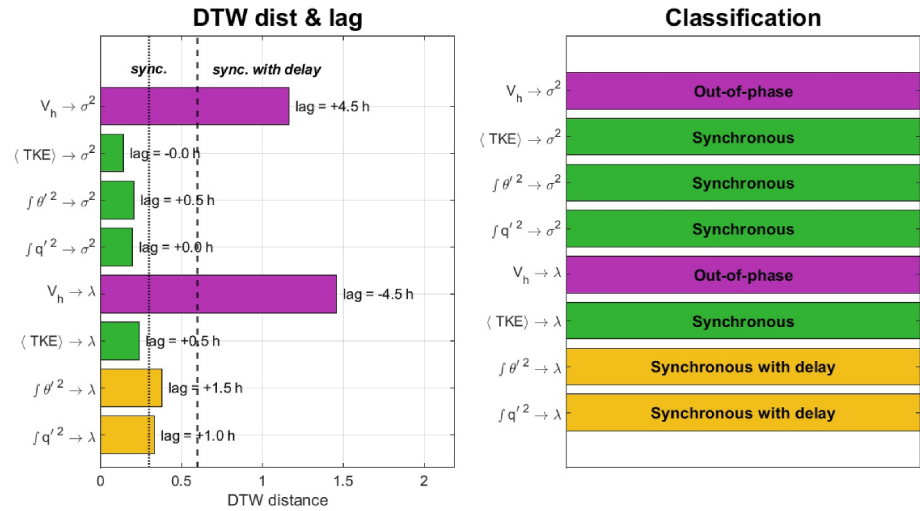


Figure 4. Diurnal cycles of large-eddy simulations predictor-target pairs after standardization. Top row: targets are λ . Bottom row: targets are σ^2 . Colored background shading indicates the four diurnal phases (Growth, Peak, Decay, Plateau). Red triangles mark principal peaks, blue triangles principal troughs. Each panel reports the Dynamic Time Warping (DTW) distance, the optimal DTW lag and the corresponding classification.

as the boundary layer stabilizes and the geostrophic wind takes over from surface friction. This nocturnal recovery is specific to the idealized configuration and is not representative of observed clear-sky conditions where synoptic forcing modulates V_h throughout the day. We restrict the physical interpretation to the convective phase, during which V_h decreases as turbulent mixing redistributes momentum and reduces near-surface shear while λ and σ^2 both rise. The high DTW distance captures this counter-evolution. The DTW lag sign, dominated by the late-evening rebound of V_h , is not over-interpreted.

The DTW classification is summarized in Table 1 together with the lags between principal extrema.

Table 1
Diurnal Coherence Between Large-Eddy Simulations Predictors and ZWD Spectral Parameters

Predictor	Target	Extrema lag (h)		DTW		Classification
		Peak	Trough	Dist	Lag (h)	
$f \overline{q'^2}$	λ	+0.5	+3.0	0.33	+1.0	Synchronous with delay
$f \overline{\theta'^2}$	λ	+0.5	-9.5	0.38	+1.5	Synchronous with delay
$\langle \text{TKE} \rangle$	λ	+1.0	+3.0	0.24	+0.5	Synchronous
V_h	λ	+8.0	-6.5	1.46	-4.5	Out of phase
$f \overline{q'^2}$	σ^2	-0.5	+0.0	0.20	+0.0	Synchronous
$f \overline{\theta'^2}$	σ^2	-0.5	-12.5	0.21	+0.5	Synchronous
$\langle \text{TKE} \rangle$	σ^2	+0.0	+0.0	0.14	+0.0	Synchronous
V_h	σ^2	+7.0	-9.5	1.17	+4.5	Out of phase

Note. Lags between principal extrema (peak and trough) quantify temporal phase relationships (negative value: predictor leads). DTW distance and optimal lag summarize the global dynamic similarity between standardized time series. The same sign convention applies to the DTW lag. Classification follows DTW criteria: synchronous if $|\Delta t_{\text{DTW}}| \leq 0.5$ h and $d_{\text{DTW}} < 0.30$, synchronous with delay if $|\Delta t_{\text{DTW}}| \leq 2$ h and $d_{\text{DTW}} < 0.60$, out of phase if $d_{\text{DTW}} \geq 0.60$.

3.3. Phase-Resolved Slope Ratios

Table 2 reports the normalized slope ratios for the eight pairs across the four diurnal phases (Growth, Peak, Decay, Plateau), capturing the local rate of covariation in addition to the global DTW similarity.

During the Growth phase, the three turbulence-related predictors show positive slope ratios against λ between 2.55 and 2.70, meaning that they reach a substantial fraction of their daily amplitude while λ is still rising more gradually within the 9–12 hr window. This asymmetry is consistent with the small lead of λ identified by the DTW analysis: λ has already accumulated part of its diurnal amplitude before 9 hr, so it traverses a smaller relative span during Growth than the predictors that are still emerging from their morning minimum. By the Decay phase, the slope ratios approach unity for $\langle \text{TKE} \rangle$ - λ (+1.22), meaning that the spectral break collapses at the same rate as the bulk turbulence intensity once the surface forcing diminishes. Slope ratios involving V_h stay negative throughout the convective day, confirming the counter-evolution between advection velocity and the spectral parameters during this period.

3.4. Summary of LES Findings

A clear hierarchy emerges. σ^2 is dynamically synchronous with the bulk turbulence intensity, whether measured by $\langle \text{TKE} \rangle$ or by the integrated scalar variances. Such synchrony reflects the direct physical chain in which mixing produces humidity and temperature fluctuations that propagate without delay

Table 2

Normalized Rates of Change per Diurnal Phase, Expressed as a Fraction of Each Variable's Diurnal Amplitude per Hour

Predictor	Target	Growth	Peak	Decay	Plateau
$\int \overline{q'^2}$	λ	+2.67	+0.73	+0.87	−0.28
$\int \overline{\theta'^2}$	λ	+2.70	+0.38	+0.80	−0.45
$\langle \text{TKE} \rangle$	λ	+2.55	+0.18	+1.22	−0.09
V_h	λ	−0.81	−0.57	−1.12	−0.19
$\int \overline{q'^2}$	σ^2	+1.45	+3.17	+0.88	+0.42
$\int \overline{\theta'^2}$	σ^2	+1.47	+1.65	+0.81	+0.67
$\langle \text{TKE} \rangle$	σ^2	+1.38	+0.76	+1.23	+0.14
V_h	σ^2	−0.44	−2.45	−1.13	+0.29

Note. The ratio $r^{(p)} = s_x^{(p)}/s_y^{(p)}$ quantifies whether the predictor evolves faster (ratio above unity), slower (between zero and unity), or in opposite direction (negative ratio) compared to the target during each phase. Phase boundaries: Growth 9–12 hr, Peak 12–15 hr, Decay 15–19 hr, Plateau 19–21.5 hr.

to the path-integrated wet refractivity. λ , in turn, is synchronous with $\langle \text{TKE} \rangle$ and synchronous with delay relative to the scalar variances, with λ leading by half an hour to one and a half hours. This lead spans several multiples of the convective eddy-turnover timescale $\tau_* = z_i/w_*$, which is of order 10–20 min for the simulated conditions. The ordering supports the reading that the spectral transition becomes detectable in the periodogram as soon as the first organized eddies emerge, while the path-integrated scalar variances accumulate over multiple eddy-turnover times before reaching their full diurnal amplitude. As for V_h , it is dynamically out of phase with both spectral parameters during the convective period, reflecting the counter-evolution of mean wind and turbulence intensity in an idealized configuration without competing synoptic forcing.

These results refine the interpretation of λ proposed in earlier studies (Kermarrec & Schön, 2020; Kermarrec, Calbet, et al., 2025; Kermarrec, Klos, et al., 2023). Under the convective regime simulated here, λ does not reduce to a Taylor frozen-flow proxy of advection velocity. Its evolution is more closely tied to the dynamic state of the turbulent cascade, with the bulk turbulence intensity and the scalar variances acting as the primary drivers and V_h acting as a counter-evolving variable rather than a direct controller. We return to this point in Section 5, where we discuss its implications for the regime-dependent meaning of λ .

4. Observational Validation

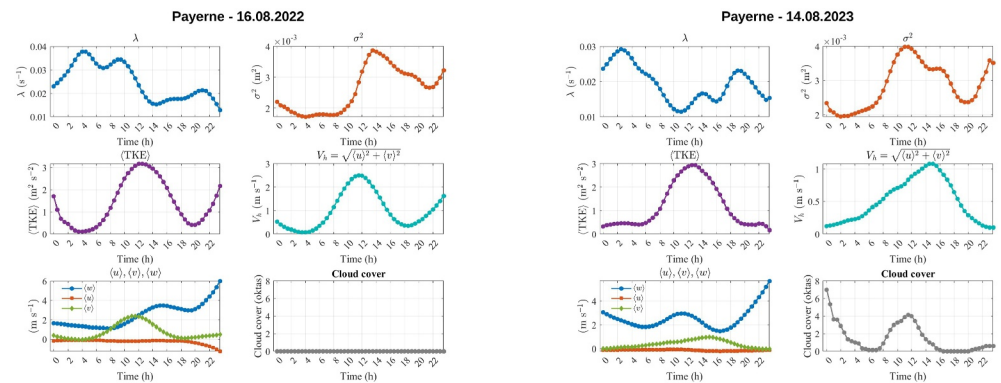
The LES analysis identifies coherence patterns under idealized conditions. We now test whether these relationships hold in real atmospheric observations, using 3 years of co-located GNSS and Doppler LiDAR measurements from Payerne. We first examine individual days through case studies, then apply statistical analysis across the full data set.

4.1. Approach

The LES analysis provides a coherence framework for the temporal coupling between boundary-layer variables and ZWD spectral parameters under idealized convective conditions. We now test whether the same framework can be transferred to real atmospheric observations, using 3 years of co-located GNSS and Doppler wind LiDAR measurements at Payerne. The transfer is non-trivial: observations are subject to synoptic forcing absent from the simulation, the LiDAR provides only $\langle \text{TKE} \rangle$ and V_h rather than the four LES predictors, and the diurnal cycles span the full 24 hr rather than the 12.5 hr convective window of the LES.

We retain the four pairs $\langle \text{TKE} \rangle - \lambda$, $V_h - \lambda$, $\langle \text{TKE} \rangle - \sigma^2$ and $V_h - \sigma^2$. The diurnal coherence framework defined in Section 2.5 is applied without methodological modification, with two adjustments justified by the longer observation window. The phase decomposition is extended to five phases, with an additional Pre-dawn phase covering 0–6 hr. The detection of principal extrema for the lag analysis is restricted to the 06–22 hr window to avoid spurious matches near midnight, and all lags are wrapped to $[-12, +12]$ h. We retain only days with mean cloud cover $\overline{\text{CC}} \leq 2$ oktas, which corresponds to predominantly clear-sky conditions while permitting transient passages of clouds, yielding 249 days out of 1,096. We stratify the analysis by season, distinguishing summer (June–August, 109 days), winter (December–February, 33 days) and the transition seasons (March–May and September–November, 105 days). The seasonal stratification serves a specific purpose: summer days at Payerne provide the closest analog to the convective regime simulated by the LES, while winter days, dominated by stable stratification under anticyclonic conditions, offer a contrasting regime against which the framework can be tested.

We organize the analysis in three parts. Section 4.2 examines four representative days that illustrate the diurnal coherence in distinct dynamical regimes. Section 4.3 presents the aggregated statistics across the 249 clear-sky days with seasonal stratification. Section 4.4 compares the simulated and observed coherence patterns on a common basis.



(a) 16 August 2022: strong advection.

(b) 14 August 2023: weak advection.

Figure 5. Summer case studies at Payerne illustrating two distinct regimes. (a) Strong advection ($V_h \in [0.5, 3] \text{ m s}^{-1}$): σ^2 tracks $\langle \text{TKE} \rangle$ synchronously, while λ and V_h are out of phase. (b) Weak advection ($V_h < 1 \text{ m s}^{-1}$): σ^2 tracks $\langle \text{TKE} \rangle$ synchronously, but λ no longer mirrors V_h and instead reflects the size of dominant convective eddies.

4.2. Case Studies

We use four selected days from the Payerne data set, each illustrating a distinct dynamical regime: strong advection in summer, weak advection in summer, spring convection at intermediate V_h , and a winter cloudy day. The cases are chosen to span the parameter space rather than to selectively support the framework, and we describe the coherence patterns in their own terms before turning to the aggregated statistics.

4.2.1. Strong Advection Summer (16 August 2022)

This day (Figure 5a) combines persistent clear skies (0 oktas) with a substantial diurnal range of V_h between 0.5 and 3 m s^{-1} . Such conditions provide a clean test of the framework. $\langle \text{TKE} \rangle$ follows a textbook convective cycle, rising from about $1.5 \text{ m}^2 \text{ s}^{-2}$ in the morning to $3.5 \text{ m}^2 \text{ s}^{-2}$ near midday and decaying to $0.5 \text{ m}^2 \text{ s}^{-2}$ by evening. The variance σ^2 tracks $\langle \text{TKE} \rangle$ in shape and timing, with a peak at $3.5 \times 10^{-3} \text{ m}^2$ in the early afternoon and a synchronous decay through the late afternoon. This pair is consistent with the synchronous classification identified in the LES.

A different pattern emerges for the λ - V_h pair. The cutoff frequency peaks at 0.035 s^{-1} during the morning when V_h is at its minimum, reaches a local minimum around 0.02 s^{-1} at midday when V_h is highest, and recovers in the late afternoon as V_h decreases. The peak-to-trough lag between λ and V_h is approximately 3 hrs and the shape mismatch is pronounced, consistent with an out-of-phase coupling. Such a signature is expected when advection variability is large enough to dominate the spectral break: increases in V_h compress the temporal signature of turbulent structures and redistribute spectral energy toward higher frequencies, which manifests as a decrease in λ .

4.2.2. Weak Advection Summer (14 August 2023)

This case (Figure 5b) presents the converse situation in which advection variability is suppressed, with V_h remaining below 1 m s^{-1} throughout the day. $\langle \text{TKE} \rangle$ still traces a clear convective cycle peaking at $3 \text{ m}^2 \text{ s}^{-2}$, and σ^2 follows it closely, again pointing to synchronous coupling. By contrast, λ no longer mirrors the anti-phased pattern of the strong-advection case. Instead, the spectral break oscillates with minima around 10 and 16 hr, both coinciding with peaks of convective activity, while its maxima fall during the morning and evening transitions. Under this regime, λ tracks the dominant eddy size in the energy-input range: large thermals during peak heating shift energy toward lower frequencies and reduce λ , while smaller and more diverse structures during transitions populate higher frequencies and increase λ . This dual behavior, in which the controlling mechanism for λ depends on whether advection variability is large or small, is the empirical motivation for the regime-dependent interpretation of λ developed in the discussion.

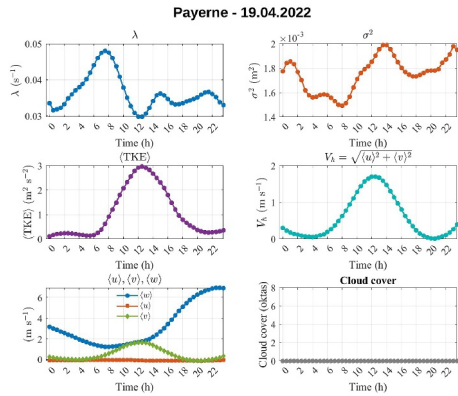


Figure 6. Spring case at Payerne (19 April 2022). Clear-sky conditions with intermediate V_h variability. The σ^2 – $\langle\text{TKE}\rangle$ pair is synchronous and the λ – V_h pair is out of phase.

4.2.3. Spring Convection (19 April 2022)

The spring case (Figure 6) features a textbook clear-sky convective cycle, with $\langle\text{TKE}\rangle$ growing from near zero to about $3 \text{ m}^2 \text{ s}^{-2}$ by early afternoon. σ^2 parallels $\langle\text{TKE}\rangle$ in shape and timing, providing a third instance of synchronous coupling outside the summer season. V_h ranges from below 0.5 m s^{-1} in the morning to about 2 m s^{-1} at midday, intermediate between the two summer cases. As in the strong-advection summer day, the cutoff frequency exhibits a pronounced morning peak at 0.048 s^{-1} when V_h is at its minimum, then decreases monotonically to about 0.03 s^{-1} as V_h rises through midday. The shape complementarity between the two series matches the out-of-phase coupling seen in summer, indicating that the framework extends to the spring season under sufficient advection variability.

4.2.4. Winter Cloudy Conditions (18 January 2024)

The winter case (Figure 7) tests the framework under conditions far from the LES setup. Cloud cover stays between 6 and eight oktas throughout the day, falling outside the clear-sky filter applied for the aggregated analysis but informative as a single-day comparison. Despite suppressed surface heating, $\langle\text{TKE}\rangle$ reaches $2.3 \text{ m}^2 \text{ s}^{-2}$, comparable to summer values. Mechanically forced or cloud-top driven mixing must therefore sustain the boundary-layer turbulence. The variance σ^2 continues to track $\langle\text{TKE}\rangle$ in shape and amplitude, indicating that the synchronous coupling between turbulence intensity and ZWD variance persists outside the convective regime.

The cutoff frequency, by contrast, shows limited diurnal variability, consistent with the absence of strong convective forcing. V_h remains nearly constant, so that the out-of-phase coupling cannot be assessed: with one of the two predictors flat, the diurnal coherence framework yields no informative result for the V_h – λ pair on this particular day. Such a non-result is itself instructive. It illustrates the regime-dependence of the λ – V_h relationship: when neither convective forcing nor advection variability is strong, λ reflects whatever residual structure the boundary layer retains, and the framework returns no clear classification rather than a spurious one.

4.3. Aggregated Diurnal Coherence

We now apply the diurnal coherence framework to the 249 clear-sky days, with the seasonal stratification described in Section 4.1.

4.3.1. Pairwise Classification by Season

Table 3 reports the median DTW distances and lags for the four pairs $\langle\text{TKE}\rangle$ – σ^2 , $\langle\text{TKE}\rangle$ – λ , V_h – σ^2 and V_h – λ , separately for the summer, winter and transition seasons. Three patterns emerge.

The first concerns the strongest coupling. Across the year, $\langle\text{TKE}\rangle$ – σ^2 shows the smallest DTW distances ($d_{\text{DTW}} = 0.45$ in summer, 0.52 in transition, 0.66 in winter), with the optimal lag stably around -2 h : $\langle\text{TKE}\rangle$ leads σ^2 by approximately 2 hrs. Both distance and lag exceed the LES values, where the pair is synchronous with a near-zero lag. The shift toward synchronous-with-delay has a clear physical reading. The LiDAR-measured TKE (100–850 m) captures bulk turbulent intensity in the developing convective layer, while σ^2 is dominated by path-integrated wet-refractivity variance, concentrated where the vertical gradient of specific humidity is largest—typically near the boundary-layer top and the entrainment zone. The LiDAR TKE peaks while bulk mixing is most vigorous, whereas σ^2 continues to grow as the boundary layer deepens through layers of contrasting humidity. The 2-hour lag reflects this offset between bulk mixing intensity and the moisture variance signature tied to the boundary-layer top trajectory.

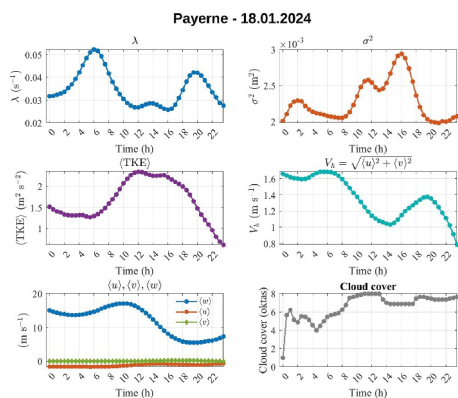


Figure 7. Winter case at Payerne (18 January 2024). Cloudy conditions (6–8 oktas) with limited V_h variability. The σ^2 – $\langle\text{TKE}\rangle$ coupling persists, while the λ – V_h coupling is not informative under flat V_h .

Table 3
Diurnal Coherence Between Payerne Predictors and ZWD Spectral Parameters, Stratified by Season

Pred	Target	Summer (<i>N</i> = 109)		Winter (<i>N</i> = 33)		Transition (<i>N</i> = 105)	
		Dist	Lag (h)	Dist	Lag (h)	Dist	Lag (h)
⟨TKE⟩	σ^2	0.45	−2.0	0.66	−1.0	0.52	−1.5
⟨TKE⟩	λ	0.61	+3.5	0.68	+0.5	0.64	+2.0
V_h	σ^2	0.60	−2.0	0.66	−1.0	0.62	−1.5
V_h	λ	0.66	+2.4	0.70	+0.5	0.68	+1.5

Note. Summer: June–August. Winter: December–February. Transition: March–May and September–November. Values are medians across days within each season. Classification criteria as in Table 1.

A second pattern emerges for $V_h-\sigma^2$. DTW distances cluster around 0.60 in summer and rise to 0.66 in winter, classifying the pair as out of phase, in agreement with the LES result. The lag stays centered around −2 h. For an anti-correlated pair, the sign of the DTW lag is sensitive to the precise alignment of secondary features and should not be over-interpreted: the relevant outcome is the persistence of the out-of-phase classification, which reflects the well-known counter-evolution of mean wind and turbulence intensity during the convective cycle. As expected, the signal weakens in winter, when the diurnal cycle of V_h flattens under stable stratification.

A third pattern concerns the pairs involving λ . ⟨TKE⟩− λ has $d_{DTW} = 0.61$ with a lag of +3.5 h, indicating that λ leads ⟨TKE⟩ by about three and a half hours. This lead echoes the strong-advection case study (Section 4.2), in which λ reaches its maximum in late morning when V_h is minimum, hours before the LiDAR TKE peaks in early afternoon; the same ordering dominates the summer aggregate. The lead should not be read as a causal precedence; it reflects a temporal ordering between two physically distinct observables of boundary-layer convection sampled by instruments with different vertical sensitivities. The $V_h-\lambda$ pair has $d_{DTW} = 0.66$ with a lag of +2.4 h and remains out of phase; as for $V_h-\sigma^2$, the DTW lag sign is not over-interpreted. In winter, both distances climb to 0.68–0.70 and the lags lose interpretability as the diurnal cycles flatten.

The seasonal stratification confirms that the framework transfers most cleanly to summer days, the closest analog of the simulated convective regime. Winter days, dominated by stable stratification and reduced diurnal forcing, do not exhibit the coherence signatures—a result that confirms the framework does not produce false positives where no convective coupling is expected.

The seasonal stratification confirms that the framework transfers most cleanly to summer days, the closest analog of the simulated convective regime. Winter days, dominated by stable stratification and reduced diurnal forcing, do not exhibit the coherence signatures—a result that confirms the framework does not produce false positives where no convective coupling is expected.

4.3.2. Significance of the Extrema Opposition Ratio

The original analysis of this study used an extrema-based ratio to quantify the coupling between λ and V_h : the fraction of days on which a peak in one series falls within a narrow temporal tolerance of a trough in the other. This local test is stricter than the global shape similarity assessed by DTW, and it warrants a formal significance test, which we report here.

4.3.2.1. Methodology

For each day, we detect peaks and troughs in the smoothed diurnal cycles of λ and V_h within the 06–22 hr window, using a prominence threshold of 25% of the daily amplitude, and we count the fraction of pairings that fall in opposition within a tolerance $\tau = 1.5$ h. The Monte Carlo permutation test described in Section 2.5 provides a one-sided *p*-value for the null hypothesis of no temporal coupling. We apply the test on two scopes: the full set of 1,074 days for which the detection succeeds, and the 249-day clear-sky subset on which the rest of the aggregated analysis operates.

4.3.2.2. Results

Across the 1,074 days, the observed opposition ratio is $\hat{r} = 0.502$ ($Z = 0.3, p = 0.38$). On the 249-day clear-sky subset, $\hat{r} = 0.510$ ($Z = 0.5, p = 0.31$). Figure 8 compares the two null distributions. In neither scope is the observed ratio statistically distinguishable from chance at the 0.05 level. Sensitivity to the detection parameters is illustrated in Figure 9: across the grid of prominence thresholds {0.15, 0.20, 0.25, 0.30, 0.35} and tolerances {0.5, 1.0, 1.5, 2.0, 2.5, 3.0} h, the ratio remains in the range [0.499, 0.557], confirming that no choice of parameters within reasonable bounds yields a significantly anti-phased signal at the level of individual extrema.

4.3.2.3. Joint Reading With the DTW Classification

The extrema test complements rather than contradicts the DTW classification of the same data set, which identifies $\lambda-V_h$ as out of phase in summer ($d_{DTW} = 0.66$, lag +2.4 h). DTW evaluates the global similarity of two diurnal cycles tolerating physical lags within its alignment search, while the extrema ratio evaluates a stricter local property: whether individual peaks and troughs fall within a narrow tolerance on each day. The two outcomes




Figure 8. Monte Carlo permutation test of the $\lambda-V_h$ extrema opposition ratio at Payerne. Histograms show the null distribution of opposition ratios from 5,000 permutations in which the times of the V_h extrema are randomized within the 06–22 hr window. Left: full 1074-day sample. Right: 249-day clear-sky subset.

coexist consistently here—the average diurnal cycles are dynamically out of phase, but the day-by-day alignment of individual extrema is too dispersed for the local test to be significant. The dispersion is itself physically meaningful, reflecting 3 years of synoptic-forcing variability. We use the DTW classification as the central evidence for the $\lambda-V_h$ out-of-phase coupling and retain the extrema-based result as a transparent complement.

4.3.3. Bootstrap Uncertainty of λ

The interpretation of λ developed in this study relies on its diurnal variation, including its low nocturnal values, so the estimation uncertainty must be quantified. We propagate the spectral fit uncertainty by bootstrap resampling of the 10-second batches contributing to each 30-min Whittle estimate, with 200 resamplings per window. Figure 10 shows the diurnal cycle of the median λ at Payerne with the bootstrap standard error overplotted. The relative uncertainty σ_λ/λ averages 2.3% during daytime hours and 3.1% at night, with a maximum around 7% at the lowest λ in the early morning. These values remain well below the diurnal amplitude of λ (about 30% of the daily mean), confirming that the nocturnal values are physically resolved and not artifacts of reduced spectral resolution.

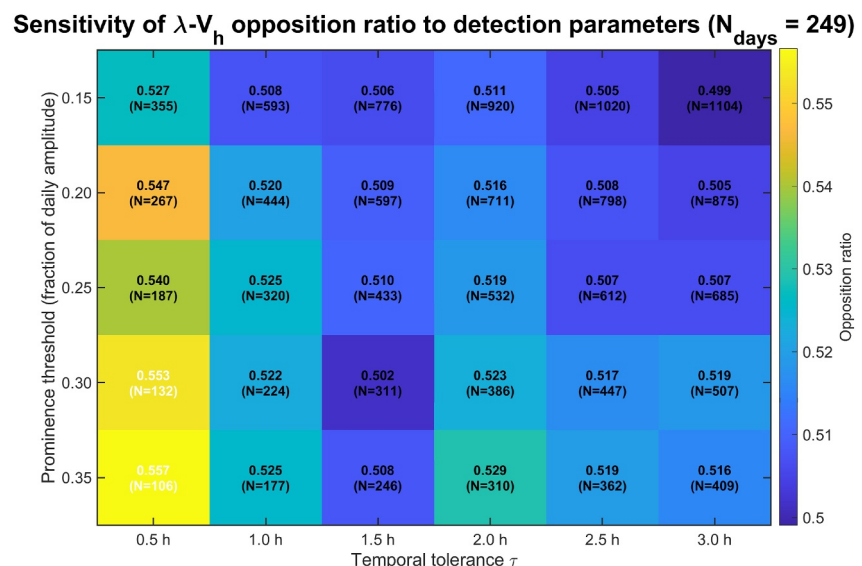


Figure 9. Sensitivity of the $\lambda-V_h$ opposition ratio to the prominence threshold (rows) and the temporal tolerance (columns), evaluated on the 249 clear-sky days. Each cell reports the ratio and the total number of extrema events N .

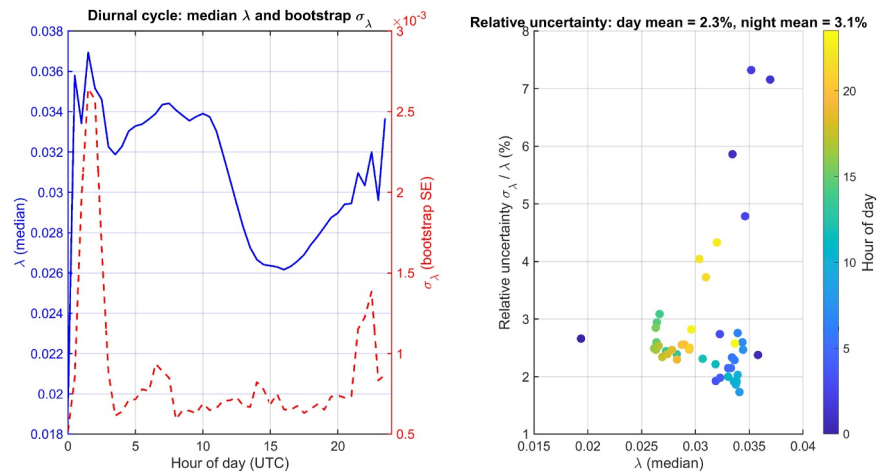


Figure 10. Bootstrap uncertainty of λ at Payerne. Left: median λ across the 249 clear-sky days (blue) and bootstrap standard error σ_λ (red dashed) as a function of the hour of day. Right: relative uncertainty σ_λ/λ as a function of λ , colored by the hour of day.

4.4. Comparison Between LES and Summer Observations

Bringing the LES and Payerne analyses together provides a transferability test of the framework. Figure 11 compares the DTW distances and classifications for the four pairs common to both data sets, restricted to the summer subset of the observations (109 days).

The qualitative classification is largely preserved. Both pairs involving V_h classify as out of phase in either data set, $\langle \text{TKE} \rangle - \sigma^2$ stays within the synchronous or synchronous-with-delay regime, and $\langle \text{TKE} \rangle - \lambda$ sits at the boundary between synchronous-with-delay and out-of-phase ($d_{\text{DTW}} = 0.24$ in LES, 0.61 in summer Payerne, with a $\pm 2\%$ threshold adjustment flipping the latter). DTW distances grow systematically from LES to observations, reflecting day-to-day variability of synoptic forcing, partial spatial overlap between the LiDAR volume and the GNSS receiver, and variable boundary-layer depth across days. Even with this absolute increase, the relative ordering of pairs is preserved, with $\langle \text{TKE} \rangle - \sigma^2$ remaining the most coherent pair in both data sets.

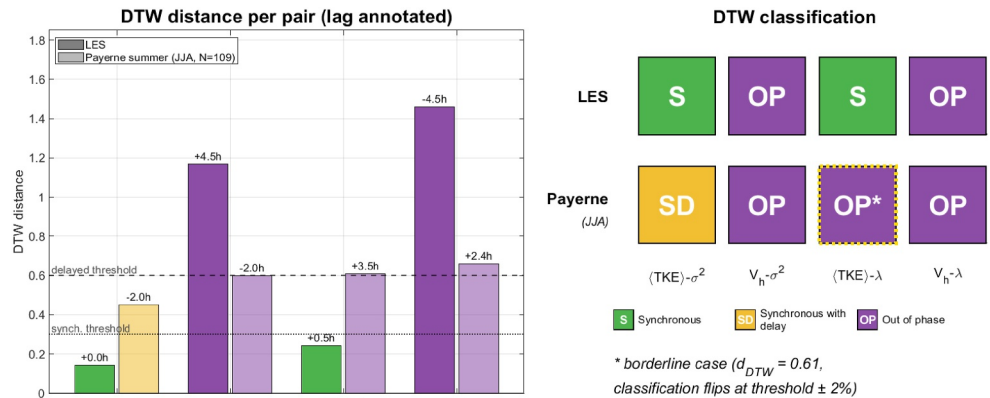


Figure 11. Comparison between large-eddy simulations (LES) (left bars, opaque) and the summer subset of the Payerne observations (right bars, semi-transparent) for the four pairs $\langle \text{TKE} \rangle - \sigma^2$, $V_h - \sigma^2$, $\langle \text{TKE} \rangle - \lambda$ and $V_h - \lambda$. Left panel: Dynamic Time Warping (DTW) distance, with the synchronous threshold (dotted line at 0.30) and the synchronous-with-delay threshold (dashed line at 0.60) marked. Each bar is annotated with the optimal DTW lag in hours. Bar color indicates classification (green: synchronous; yellow: synchronous with delay; purple: out of phase). Right panel: synoptic table of classifications with short codes (S, SD, OP), LES on top and Payerne summer on bottom. The dotted yellow border on Payerne $\langle \text{TKE} \rangle - \lambda$ marks a borderline case ($d_{\text{DTW}} = 0.61$ vs. the synchronous-with-delay threshold at 0.60).

Table 4
Summary of LES–Observation Agreement Across Case Studies and Aggregated Analyses

Pair	Case studies				Summer Aggregated	Winter Aggregated
	16.08.22	14.08.23	19.04.22	18.01.24		
σ^2 tracks $\langle \text{TKE} \rangle$	Y	Y	Y	Y	Y	P
V_h – σ^2 out of phase	Y	N	Y	N	Y	N
$\langle \text{TKE} \rangle$ – λ coupled with lag	Y	Y	Y	N	Y	N
V_h – λ out of phase	Y	N	Y	N	Y	N

Note. Y: clear coherence consistent with LES; P: partial coherence (one cycle too flat to be informative); N: regime in which the LES classification does not transfer.

The lags differ in an informative way. In the LES, $\langle \text{TKE} \rangle$ and σ^2 are synchronous and λ leads $\langle \text{TKE} \rangle$ by half an hour. In summer Payerne, $\langle \text{TKE} \rangle$ leads σ^2 by about 2 hrs and λ leads $\langle \text{TKE} \rangle$ by about three and a half hours. These longer lags are not unexpected: the LES is an idealized single-cycle simulation, while the Payerne medians aggregate 3 years of clear-sky days with day-to-day variability of synoptic forcing, surface heterogeneity, boundary-layer depth and instrumental geometry. None of these factors is captured by the simulation, and disentangling their individual contributions is beyond the scope of the present study. The relevant point for transferability is that the qualitative classification is preserved across the two data sets despite this quantitative widening. Table 4 summarizes the agreement across case studies and aggregated analyses.

5. Discussion

5.1. Regime-Dependent Interpretation of λ

The variance σ^2 has a straightforward interpretation as the path-integrated variance of the wet refractivity, driven by humidity and temperature fluctuations through the Smith and Weintraub formulation given in Equation 1. The synchronous coupling identified by the DTW analysis confirms that σ^2 is dominated by bulk turbulence intensity. In the LES, where all variables share the same vertical extent, the phase shift relative to $\langle \text{TKE} \rangle$ is negligible; in the Payerne observations the coupling is preserved but offset by approximately 2 hrs, in line with the broader widening of lags between LES and observations discussed in Section 4.4 and not specific to a particular physical regime.

The interpretation of λ is more subtle. As the cutoff frequency, λ marks the temporal scale at which the ZWD spectrum transitions from the energy-input range to the inertial subrange. Under Taylor's frozen-flow hypothesis, this scale relates to the advection velocity and the outer scale of turbulence: faster advection compresses the temporal signature of structures of a given size, shifting λ toward higher frequencies. Such an advective interpretation would predict a positive correlation between λ and V_h .

Yet both the LES and the summer Payerne observations yield an out-of-phase coupling instead. The strong-advection summer day (16 August 2022) provides the cleanest illustration: λ peaks when V_h reaches its minimum, drops to a local minimum when V_h peaks, and recovers as V_h decreases. Resolving this apparent paradox requires invoking the dynamical coupling between advection velocity and turbulence intensity within the convective boundary layer. During the development phase, vigorous mixing redistributes momentum across the column and reduces near-surface shear, lowering V_h while populating the boundary layer with rapidly evolving structures of decreasing size. The resulting spectrum has its energy-input range pushed toward higher frequencies, yielding higher λ . After the convective peak, the boundary layer stabilizes, V_h recovers, and turbulent structures collapse into longer-lived horizontally elongated patterns at lower λ . The two variables are anti-correlated over the diurnal cycle not because advection has no effect on the spectrum, but because advection variability and turbulence intensity are themselves coupled through the boundary-layer mass and momentum budgets.

The weak-advection summer case (14 August 2023) confirms this interpretation by complementarity. With V_h below 1 m s^{−1} throughout the day, the advective contribution to λ becomes essentially constant, and the diurnal variability of λ is driven instead by the size of dominant convective structures. Large thermals during peak heating dominate the low-frequency end of the spectrum, lowering λ , while smaller structures during transitions populate

higher frequencies and increase λ . The cutoff frequency therefore carries different physical information depending on the dynamical regime: it tracks the advection-modulated temporal compression of structures when V_h varies substantially, and the dominant eddy size when V_h is approximately constant. Such regime-dependence is intrinsic to the spectral parameter rather than an artifact of the analysis. It complicates the use of λ as a single-purpose proxy and suggests that any interpretation of λ should be informed by the concurrent state of V_h .

5.2. Other Sources of ZWD Variability and the Selectivity of the Von Kármán Fit

ZWD fluctuations on the timescales considered here are not driven exclusively by boundary-layer turbulence. Convective events and intense precipitation produce non-stationary refractivity excursions (Ma et al., 2022; Nahmani et al., 2019), and high-frequency components of the gravity-wave field can contribute spectral energy in the same band (Fritts & Alexander, 2003). The present analysis is restricted to the turbulent component by three complementary filtering mechanisms.

The parametric von Kármán fit captures only spectral shapes consistent with a Kolmogorov-type cascade with a finite outer scale, so that gravity-wave signatures, which manifest as resonant peaks or extended plateaus rather than the smooth plateau-to-cascade transition, are rejected at the fit stage. The inertia-gravity wave continuum is in any case centered at frequencies below 10^{-3} Hz, well below our analysis band of 5×10^{-4} to 5×10^{-2} Hz. The lower bound is set by the 30-min window length and the upper bound is conservatively kept an order of magnitude below the Nyquist frequency to avoid contamination by GNSS receiver noise and unmodelled high-frequency multipath, which dominate the periodogram above approximately 0.1 Hz. Convective transients introduce a separate failure mode: by violating the second-order stationarity assumption of the Whittle estimator, they produce fitting residuals that exceed our 3σ rejection threshold and are likewise filtered out at the fit stage. As a coarser-grained protection, the day-level clear-sky filter ($\overline{CC} \leq 2$ oktas) excludes days dominated by deep convection from the aggregated analysis. Combined, these mechanisms restrict the framework to fluctuations consistent with boundary-layer turbulence by construction. We acknowledge that residual contamination cannot be excluded, and that more granular regime classifications, for instance based on synoptic typing or independent gravity-wave proxies, would refine the present approach.

5.3. Implications and Limitations

The synchronous coupling between σ^2 and bulk turbulence intensity supports the use of ZWD variance as a column-integrated proxy of mixing under convective conditions, complementary to point measurements from towers or sparse radiosonde profiles. With appropriate calibration against co-located reference instruments, σ^2 retrieved from existing GNSS networks could provide continuous monitoring of boundary-layer mixing intensity in regions lacking dedicated instrumentation. As for λ , it is most informative when interpreted jointly with σ^2 and concurrent V_h . The advantage of GNSS lies in spatial coverage: more than 30,000 stations operate globally (International GNSS Service, 2023) with continuous temporal sampling, complementing rather than replacing existing remote-sensing infrastructure.

Several limitations deserve mention. The LES represents a single diurnal cycle under idealized conditions, and the geographic generalization rests on a single mid-latitude site at Payerne. The forthcoming FESSTVaL intensive observation period in the Berlin-Potsdam region (Hohenegger et al., 2023) will provide a multi-station deployment of GNSS receivers and Doppler wind LiDARs under contrasting boundary-layer regimes, including urban-rural transitions and convergence lines associated with mesoscale heterogeneity. A second limitation concerns the predictors retrievable from the Doppler LiDAR: integrated humidity and temperature variances are not available from this instrument, restricting the aggregated analysis to the four pairs that involve $\langle \text{TKE} \rangle$ and V_h . Combining GNSS with co-located Raman LiDARs or differential absorption LiDARs would close this gap. A third direction concerns the extension of the framework to the pre-convective regime, where the spectral signatures of boundary-layer destabilization are of operational interest for nowcasting and where adaptation of the framework to non-clear-sky conditions would be required. Finally, operational use of ZWD spectral parameters requires absolute calibration of σ^2 and λ retrievals across receivers and processing pipelines, beyond the scope of the present study.

6. Conclusions

We tested the hypothesis that the spectral parameters of high-frequency ZWD fluctuations carry diagnostic information about the convective boundary-layer state. To this end, we developed a diurnal coherence framework based on DTW, extrema lags and phase-resolved slope ratios, calibrated on an idealized LES of a single convective diurnal cycle and transferred without modification to 3 years of co-located GNSS and Doppler wind LiDAR observations at Payerne.

In the LES, σ^2 is synchronous with bulk turbulence intensity, λ leads $\langle \text{TKE} \rangle$ by half an hour and trails the integrated scalar variances by 1–1.5 hr, and V_h is dynamically out of phase with both spectral parameters during the convective phase. The transfer to the Payerne observations succeeds for the summer subset, where the four pairs that involve $\langle \text{TKE} \rangle$ and V_h retain the qualitative classification of the LES. Lags between the spectral parameters and the boundary-layer variables are longer in the observations than in the LES, reflecting sources of variability that the idealized single-cycle simulation does not capture. By contrast, the framework yields no informative classification for the winter subset, where the diurnal cycles flatten — an instructive negative result that delimits the regime of applicability.

Bootstrap propagation of the spectral fit uncertainty yields relative uncertainties on λ of 2.3% during daytime and 3.1% at night, confirming that nocturnal values are physically resolved. The original extrema-based opposition ratio for the λ – V_h pair, recomputed and tested by Monte Carlo permutation, is not statistically distinguishable from chance on the clear-sky subset ($\hat{r} = 0.510$, $p = 0.31$), with a stability range across detection parameters of [0.499, 0.557]. Far from contradicting each other, the DTW classification and the extrema ratio offer complementary readings: the global shape coherence of the λ – V_h cycles is captured by the DTW distance, while the dispersion of synoptic forcing across 3 years prevents a tight day-by-day alignment of individual extrema.

The cutoff frequency λ is not a single-purpose proxy of any one boundary-layer variable, but a regime-dependent quantity that tracks the temporal compression of structures by advection when V_h varies substantially, and the size of dominant convective eddies when V_h is approximately constant. The variance σ^2 , by contrast, has a more direct interpretation as a column-integrated proxy of turbulence intensity that holds across the regimes considered here. ZWD spectral parameters thus represent an underexploited resource for boundary-layer monitoring on the existing GNSS infrastructure. Future work will extend the framework to multi-station deployments such as the FESSTVaL intensive observation period and to the pre-convective regime relevant for nowcasting.

Appendix A: Why Pearson Correlation Is Unfit for Diurnally Forced Systems

The argument that the Pearson coefficient is inappropriate as a coherence diagnostic for diurnally forced variables can be made precise. For two zero-mean stationary series $x(t)$ and $y(t)$ related by a pure time shift $y(t) = x(t - \tau)$, the Pearson coefficient is

$$r(x, y) = \frac{\text{cov}(x(t), x(t - \tau))}{\sigma_x^2} = \rho_x(\tau), \quad (\text{A1})$$

where ρ_x is the autocorrelation function of x . For a diurnal cycle dominated by its first harmonic at period $T = 24$ h, $\rho_x(\tau) \approx \cos(2\pi\tau/T)$. A physical lag of $\tau = T/4 = 6$ h yields $r = 0$ identically, and a lag of $\tau = T/2 = 12$ h yields $r = -1$ identically, regardless of the strength of the underlying coupling. The lags between $\langle \text{TKE} \rangle$ and λ identified by the LES analysis (about 0.5–1.5 hr) reduce the coefficient by 10%–30% on these grounds alone, and the multi-hour lags suggested by the Payerne observations reduce it further. The Pearson coefficient is therefore unable to discriminate weak coupling at zero lag from strong coupling with a physically meaningful delay, which is the discrimination required in this study. The DTW distance, by contrast, satisfies $d_{\text{DTW}}(x, x(\cdot - \tau)) = 0$ for any lag within the alignment window, with the lag itself reported separately as Δt_{DTW} , and it operates on standardised series so that amplitude differences across days are absorbed by the normalization.

Conflict of Interest

The authors declare no conflicts of interest relevant to this study.

Availability Statement

The PALM model is available at <https://palm-model.org> (Maronga et al., 2020). The co-located GNSS ZWD, Doppler LiDAR, and meteorological observations from Payerne used in this study are available at Kermarrec, Deng, and Hervo (2025), together with the MATLAB analysis code that implements the diurnal coherence framework (computation of DTW distances and lags, extrema detection and pairing, Monte Carlo permutation testing, sensitivity grids, and bootstrap uncertainty propagation). The repository also includes the LES output extracted from the PALM simulation and the scripts used to generate all figures of this manuscript.

Acknowledgments

This study is supported by the Deutsche Forschungsgemeinschaft under the project LES2GNSS KE 2453/4-1. LES simulations were performed on the HLRN supercomputer at NHR@ZIB. We thank MeteoSwiss for providing Payerne observations. Open Access funding enabled and organized by Projekt DEAL.

References

- Berndt, D. J., & Clifford, J. (1994). Using dynamic time warping to find patterns in time series. In *Workshop on knowledge discovery in databases (AAAI)* (pp. 359–370).
- Bevis, M., Businger, S., Herring, T. A., Rocken, C., Anthes, R. A., & Ware, R. H. (1992). Gps meteorology: Remote sensing of atmospheric water vapor using the global positioning system. *Journal of Geophysical Research*, 97(D14), 15787–15801. <https://doi.org/10.1029/92jd01517>
- Brocard, E., Philippon, R., Haefele, A., Romanens, G., Mueller, A., Ruffieux, D., et al. (2013). Raman LiDAR for Meteorological Observations, RALMO – Part 2: Validation of water vapor measurements. *Atmospheric Measurement Techniques*, 6(5), 1347–1358. <https://doi.org/10.5194/amt-6-1347-2013>
- Clough, S. A., Shephard, M. W., Mlawer, E. J., Delamere, J. S., Iacono, M. J., Cady-Pereira, K., et al. (2005). Atmospheric radiative transfer modeling: A summary of the aer codes, short communication. *Journal of Quantitative Spectroscopy & Radiative Transfer*, 91(2), 233–244. <https://doi.org/10.1016/j.jqsrt.2004.05.058>
- Crocetti, L., Scharfner, M., Zus, F., Zhang, W., Moeller, G., Navarro, V., et al. (2024). Global, spatially explicit modelling of zenith wet delay with xgboost. *Journal of Geodesy*, 98(23), 23. <https://doi.org/10.1007/s00190-024-01829-2>
- Duan, J., Bevis, M., Fang, P., Bock, Y., Chiswell, S., Businger, S., et al. (1996). Gps meteorology: Direct estimation of the absolute value of precipitable water. *Journal of Applied Meteorology and Climatology*, 35(6), 830–838. [https://doi.org/10.1175/1520-0450\(1996\)035<0830:GMDEOT>2.0.CO;2](https://doi.org/10.1175/1520-0450(1996)035<0830:GMDEOT>2.0.CO;2)
- Foken, T. (2008). *Micrometeorology*. Springer. <https://doi.org/10.1007/978-3-540-74666-9>
- Fritts, D. C., & Alexander, M. J. (2003). Gravity wave dynamics and effects in the middle atmosphere. *Reviews of Geophysics*, 41(1), 1003. <https://doi.org/10.1029/2001RG000106>
- Guerova, G., Jones, J., Douša, J., Dick, G., de Haan, S., Pottiaux, E., et al. (2016). Review of the state of the art and future prospects of the ground-based GNSS meteorology in Europe. *Atmospheric Measurement Techniques*, 9(11), 5385–5406. <https://doi.org/10.5194/amt-9-5385-2016>
- Hohenegger, C., Ament, F., Beyrich, F., Löhnert, U., Rust, H., Bange, J., et al. (2023). FESSTVaL: The field experiment on submesoscale spatio-temporal variability in Lindenberg. *Bulletin of the American Meteorological Society*, 104(10), E1875–E1892. <https://doi.org/10.1175/BAMS-D-21-0330.1>
- International GNSS Service. (2023). Igs network. Retrieved from <https://igs.org/network>
- Jones, J., Guerova, G., Douša, J., Dick, G., de Haan, S., Pottiaux, E., & van Malderen, R. (2020). Advanced GNSS tropospheric products for monitoring severe weather events and climate (COST Action ES1206 final report). <https://doi.org/10.1007/978-3-030-13901-8>
- Kermarrec, G., Calbet, X., Deng, Z., & Carbajal Henken, C. (2025). Measurement report: Can zenith wet delay from GNSS “see” atmospheric turbulence? Insights from case studies across diverse climate zones. *Atmospheric Chemistry and Physics*, 25(6), 3567–3581. <https://doi.org/10.5194/acp-25-3567-2025>
- Kermarrec, G., Czerwonka-Schröder, D., & Holst, C. (2023). Does atmospheric turbulence affect long-range terrestrial laser scanner observations? A case study in alpine region. In K. Stein & S. Gladysz (Eds.), *Environmental effects on light propagation and adaptive systems vi* (Vol. 12731). <https://doi.org/10.1117/12.2675958>
- Kermarrec, G., Deng, Z., & Hervo, M. (2025). ZWD_TKE_CC_Wv_Payerne [Collection]. *LUIS*. <https://doi.org/10.25835/do5g6qtc>
- Kermarrec, G., Klos, A., Lenczuk, A., & Bogusz, J. (2023). 01. Long-term temporal-scales of hydrosphere changes observed by gps over Europe: A comparison with grace and enso. *IEEE Geoscience and Remote Sensing Letters*, 1-1. <https://doi.org/10.1109/LGRS.2023.3345540>
- Kermarrec, G., & Schön, S. (2020). On the determination of the atmospheric outer scale length of turbulence using gps phase difference observations: The seewinkel network. *Earth Planets and Space*, 72, 1–16. <https://doi.org/10.1186/s40623-020-01308-w>
- Kumar, U., Legendre, C. P., Lee, J.-C., Zhao, L., & Chao, B. F. (2022). On analyzing GNSS displacement field variability of Taiwan: Hierarchical agglomerative clustering based on dynamic time warping technique. *Computers & Geosciences*, 169, 105243. <https://doi.org/10.1016/j.cageo.2022.105243>
- Lilly, J. M., Sykulski, A. M., Early, J. J., & Olhede, S. C. (2017). Fractional brownian motion, the matérn process, and stochastic modeling of turbulent dispersion. *Nonlinear Processes in Geophysics*, 24(3), 481–514. <https://doi.org/10.5194/npg-24-481-2017>
- Ma, Y., Liu, T., Chen, P., Yan, X., Lv, H., Jiang, W., & Wang, J. (2022). Investigation of precipitation event with combined use of GPS PWV and weather data. *IEEE Geoscience and Remote Sensing Letters*, 19, 1–5. <https://doi.org/10.1109/LGRS.2022.3174418>
- Maronga, B., Banzhaf, S., Burmeister, C., Esch, T., Forkel, R., Fröhlich, D., et al. (2020). Overview of the palm model system 6.0. *Geoscientific Model Development*, 13(3), 1335–1372. <https://doi.org/10.5194/gmd-13-1335-2020>
- Martucci, G., Matthias, V., & Bösenberg, J. (2010). Comparison of aerosol LiDAR retrieval methods at the EARLINET. *Atmospheric Chemistry and Physics*, 10, 7439–7457.
- Nahmani, S., Bock, O., & Guichard, F. (2019). Sensitivity of GPS tropospheric estimates to mesoscale convective systems in West Africa. *Atmospheric Chemistry and Physics*, 19(14), 9541–9561. <https://doi.org/10.5194/acp-19-9541-2019>
- Pearson, G., Davies, F., & Collier, C. (2009). An analysis of the performance of the ufam pulsed doppler LiDAR for observing the boundary layer. *Journal of Atmospheric and Oceanic Technology*, 26(2), 240–250. <https://doi.org/10.1175/2008jtecha128.1>
- Qiu, C., Wang, X., Li, Z., Zhang, S., Li, H., Zhang, J., & Yuan, H. (2020). The performance of different mapping functions and gradient models in the determination of slant tropospheric delay. *Remote Sensing*, 12(1), 130. <https://doi.org/10.3390/rs12010130>
- Rahives, C., Beyrich, F., & Raasch, S. (2022). Scan strategies for wind profiling with Doppler LiDAR – An large-eddy simulation (LES)-based evaluation. *Atmospheric Measurement Techniques*, 15(9), 2839–2856. <https://doi.org/10.5194/amt-15-2839-2022>
- Seidel, D. J., Ao, C. O., & Li, K. (2010). Estimating climatological planetary boundary layer heights from radiosonde observations: Comparison of methods and uncertainty analysis. *Journal of Geophysical Research*, 115(D16), D16113. <https://doi.org/10.1029/2009jd013680>

- Stull, R. B. (1988). *An introduction to boundary layer meteorology*. Springer Netherlands. <https://doi.org/10.1007/978-94-009-3027-8>
- Sykulski, A. M., Olhede, S. C., Guillaumin, A. P., Lilly, J. M., & Early, J. J. (2019). The debiased Whittle likelihood. *Biometrika*, *106*(2), 251–266. <https://doi.org/10.1093/biomet/asy071>
- Tatarski, V. I., Silverman, R. A., & Chako, N. (1961). *Wave propagation in a turbulent medium* (Vol. 18). McGraw-Hill.
- van Nes, E. H., Scheffer, M., Brovkin, V., Lenton, T. M., Ye, H., Deyle, E., & Sugihara, G. (2015). Causal feedbacks in climate change. *Nature Climate Change*, *5*, 445–448. <https://doi.org/10.1038/nclimate2568>
- Wacker, S., Gröbner, J., Zysset, C., Diener, L., Tzoumanikas, P., Kazantzidis, A., et al. (2015). Cloud observations in Switzerland using hemispherical sky cameras. *Journal of Geophysical Research: Atmospheres*, *120*(2), 695–707. <https://doi.org/10.1002/2014JD022643>
- Wheeler, A. D. (2001). *Electromagnetic scintillation* (Vol. 1). Cambridge University Press. <https://doi.org/10.1017/CBO9780511534805>
- Zus, F., Balidakis, K., Dogan, A. H., Thundathil, R., Dick, G., & Wickert, J. (2025). DNS (v1.0): An open-source ray-tracing tool for space geodetic techniques. *Geoscientific Model Development*, *18*(15), 4951–4964. <https://doi.org/10.5194/gmd-18-4951-2025>