Thermal Predictors of Open-Loop Focus Drift at the Vera C. Rubin Observatory

Christine Corry
University of California, Berkeley
Department of Astronomy
Honors Thesis

May 14, 2026

The Vera C. Rubin Observatory's Active Optics System corrects slowly varying, thermally induced defocus through closed-loop active control and, ultimately, an open-loop look-up table calibrated against bulk temperature. However, rapid residual focus drifts, quantified by the Zernike coefficient Z_4, are observed after LUT correction on timescales of $\sim$ 30 minutes. These drifts are suspected to arise from rapidly evolving thermo-optical coupling, motivating a systematic characterization of residual thermal effects on focus stability. We analyze 23 open-loop stability image sequences from November and December 2025, each consisting of 40 exposures acquired over $\sim$28 minutes with the hexapod held fixed. Temperature data from 13 sensors spanning the enclosure air, telescope structure, mirror glass, and hexapod assemblies are queried and mean-centered to isolate thermal dynamics from bulk nightly cooling. We compute Pearson correlations between Z_4 and individual sensors, physically motivated pairwise temperature gradient proxies, and PCA-derived thermal modes, repeating all analyses at the slope timescale. Individual sensor and temperature gradient proxy correlations are weak (max $r \simeq 0.3$) and no sensor maintains a stable multivariate coefficient across runs. Top sensor thermal correlations are stronger in runs classified as smooth, consistent with a thermal signal that is present but frequently obscured by measurement noise and other dynamical effects. Correlations among the Z_4 measurements themselves further complicate the interpretation, indicating that apparent sensor relationships may partly reflect shared temporal structure rather than direct causal coupling. These results suggest that thermal-optical coupling at the Rubin Observatory is distributed across multiple mechanisms and timescales and cannot be captured by simple linear relationships with individual temperature sensors. Future work will require additional temperature sensing, a larger sample of open-loop runs, and nonlinear or multivariate dynamical models to separate true thermal drivers from correlated focus variability.

1. INTRODUCTION

The Vera C. Rubin Observatory is designed to undertake an ambitious 10-year wide-field optical survey of the southern sky from Cerro Pachón in northern Chile. The project is called the Legacy Survey of Space and Time, or LSST. The Rubin Observatory is a paradigm shift in ground-based astronomy from individual object astronomy, in which researcher point telescopes at known targets, to survey-based, data-driven cosmology. Its centerpiece, the Simonyi Survey Telescope, has a three-mirror design with an 8.4-m primary/tertiary mirror (M1M3) composed of two differently curved mirrors on one borosilicate glass substrate and a 3.5-m secondary mirror (M2). This design allows high sensitivity across an exceptionally wide 3.5 degree field of view. A single image covers an area of the sky more than 40 times the size of the full moon and is taken with the largest digital camera ever made, the 3200-megapixel LSST Camera (LSSTCam). Historically, ground-based telescopes were spatially limited. To observe faint objects, they would need to hold at a single point in the sky to gather enough light to detect the object. The Rubin Observatory's sensitivity (24.5 magnitude in a 30 second exposure and 27.5 magnitude for co-added images Ivezić et al. (2019)) allows faster exposures of faint and distant objects. This combination of sensitivity, size, and fast-movement allow observation of the sky broadly over time, empowering the identification of transient events like supernovae, asteroids, and variable stars and investigation of galaxy evolution which can help to constrain the nature of cosmological features like Dark Matter and Dark Energy.

The sensitivity and resolution of the camera require precise image quality (IQ) characterized by the sharpness of the point-spread function (PSF). The driving requirements document for the Vera C. Rubin Observatory LSST survey (Ivezić et al., 2019) specifies a median delivered image quality of approximately 0.8", assuming a system contribution of 0.4" combined in quadrature with an atmospheric seeing contribution of 0.65". In

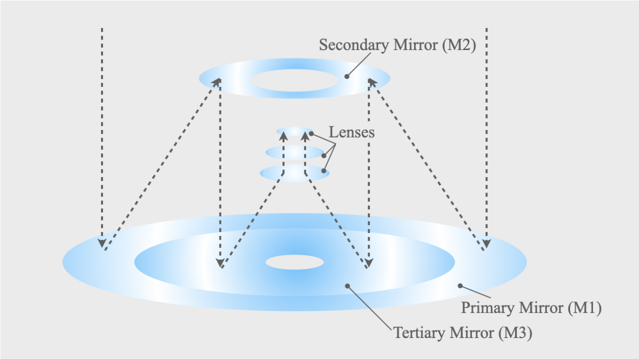

Figure 1: Schematic diagram of the Simonyi Survey Telescope optical design, showing the three-mirror system (M1, M2, M3) and refractive corrector lenses. Incident light reflects off the primary mirror (M1) to the secondary (M2), then to the tertiary (M3), before passing through the camera lenses to the focal plane. Adapted from Ivezić et al. (2019).

practice, however, the quoted 0.65" largely reflects the free-atmosphere seeing, while the total atmospheric contribution — including dome and local seeing — is typically closer to 0.8". This results in a more realistic median delivered image quality of roughly 0.95", consistent with on-sky observations.

The ideal focus behavior of an optical system would be to focus light from a point source, say a distant star, to converge on a perfect point in the image plane. Optical aberrations are deviations of the real optical system from this ideal behavior. Minimizing the point-spread corresponds to focusing light as close to a point as possible. Optimized seeing through PSF minimization can be achieved by controlling both atmospheric seeing effects and telescope-induced optical aberrations from gravitational deformation, thermal expansion, or structural misalignment. This thesis is concerned with characterizing thermally induced wavefront defocus and investigating which of components of the thermal environment within the telescope and dome are responsible for driving it.

1.1. Active Optics System

Throughout observation periods, gravitational effects, temperature, wind, and other factors can disrupt the telescope's focus, particularly as the telescope moves and tilts to image different parts of the sky. To maintain ideal image quality, the Rubin Observatory uses the Active Optics System (AOS) to dynamically orient and adjust the shape of the mirrors and correct image distortion in real-time (once every 30 s).

The AOS continuously adjusts the shapes of all three mirrors and the position of the camera and M2 relative to M1M3 with a combination of open-loop (preset) and closed-loop (real-time) controls.

The closed-loop controls are informed by corner wavefront sensors on the edges of the focal plane and address unpredictable or slowly varying aberrations. The open-loop controls operate with Look-Up Tables (LUTs) to correct predictable distortions from gravity, bulk temperature shifts, azimuth angle, and camera rotator angle. The LUT maps telescope state variables to hexapod and mirror actuator corrections to maintain optical alignment. Both are applied at the start of an exposure. The open-loop corrections are based on the current state bulk temperature and position information while the closed-loop corrections are derived from the previous image to estimate wavefront and derive corrections.

A limitation of the current LUT design is that it only corrects for the overall thermal level of the telescope, not for dynamic thermal effects or gradients between structural components. Identifying consistent dynamic thermal effects on focus would allow for more nuanced corrections in the Look-Up Tables, which would improve the convergence of the dynamic closed-loop corrections and improve the delivered image quality.

Optical Aberrations Optical aberrations are deviations from an ideal optical system, where light from a point source would be brought to a perfect focus at a single point on the focal plane. In more general terms, the optical image is not faithful to the object we are attempting to resolve. Aberrations, or wavefront errors, are complex two-dimensional functions across the pupil plane that vary with position across the focal plane. Although many decompositions can be used to describe aberrations, the choice of basis is important: different components of the aberration may have different physical causes and decomposition informs correcting these errors.

One option is Fourier modes, decomposing the wavefront into sinusoidal components. This is natural for periodic structures and computationally fast, but less physically interpretable and not well matched to a circular aperture. They are often used to characterize atmospheric turbulence (Poyneer et al., 2003). Another is Seidel aberrations which describe spherical aberration, coma, astigmatism, field curvature, and distortion in terms of the pupil and field coordinates. They have clear physical connections to the geometry of the optical system but are not orthogonal over the pupil (Welford, 1986).

For the Rubin, a particularly convenient set of decompositions is the annular Zernike polynomials.

Annular Zernike Decomposition At the Rubin Observatory, optical aberrations are characterized using annular Zernike polynomial decomposition. The wavefront error is decomposed into a set of basis functions covering Zernike coefficients from Z_4 to Z_{22}. Each coefficient describes the amplitude of a specific aberration mode.

The circular Zernike polynomials in the Noll convention Noll (1976) are defined over the unit disk as

$$Z_j(\rho,\theta) = \begin{cases} \sqrt{n+1}\, R_n^0(\rho) & m = 0 \\ \sqrt{2(n+1)}\, R_n^{|m|}(\rho)\, \cos(m\theta) & m > 0 \\ \sqrt{2(n+1)}\, R_n^{|m|}(\rho)\, \sin(|m|\theta) & m < 0 \end{cases} \tag{1}$$

where j is the Noll index, n and m are the radial and azimuthal orders respectively, $\rho \in [0,1]$ is the normalized pupil radius, θ is the azimuthal angle, and the radial polynomial is

$$R_n^{|m|}(\rho) = \sum_{s=0}^{(n-|m|)/2} \frac{(-1)^s\,(n-s)!}{s!\,\left(\frac{n+|m|}{2}-s\right)!\,\left(\frac{n-|m|}{2}-s\right)!}\,\rho^{n-2s} \tag{2}$$

The polynomials are orthonormal over the unit disk,

$$\frac{1}{\pi}\int_0^1\int_0^{2\pi} Z_j(\rho,\theta)\, Z_{j'}(\rho,\theta)\,\rho\,\mathrm{d}\rho\,\mathrm{d}\theta = \delta_{jj'} \tag{3}$$

where $\delta_{jj'}$ is the Kronecker delta. Any wavefront $W(\rho,\theta)$ can then be decomposed as

$$W(\rho,\theta) = \sum_{j=1}^{\infty} a_j\, Z_j(\rho,\theta) \tag{4}$$

where a_j are the Zernike coefficients.

The telescope's large central obscuration causes standard circular Zernike polynomials to lose their orthogonality and create a degeneracy issue. Instead, the natural basis to describe the system's aberrations are the annular Zernike polynomials because the pupil (M1) is an annulus and that is where we describe the wavefront (Xin et al., 2015).

The defocus term a_4 (corresponding to Z_4, with $n = 2$, $m = 0$) is the focus of this analysis, as it directly characterizes the axial displacement of the focal plane from its nominal position. Positive Z_4 is intra-focal (the image is in front of the nominal focus) and negative Z_4 is extra-focal (the image is behind the nominal focus).

1.2. Thermal Seeing Prior Work

Degradation of astronomical image quality by temperature gradients within and around telescope enclosures,

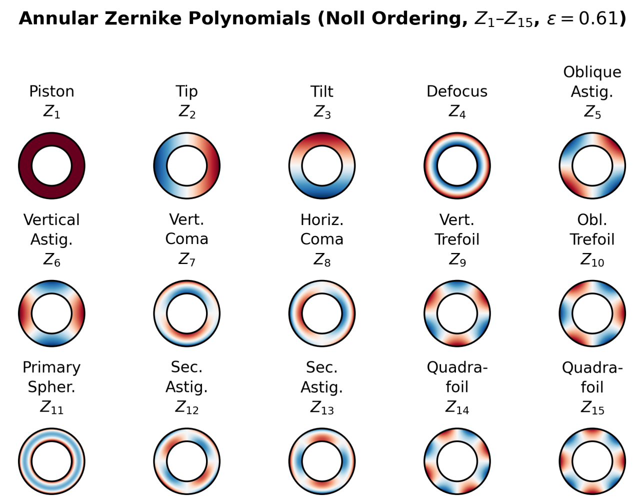

Figure 2: Annular Zernike polynomials evaluated over the Rubin pupil ($\epsilon = 0.612$; Xin et al. 2015). Red and blue lobes indicate positive and negative wavefront contributions respectively. Ordering convention follows Noll (1976).

collectively termed dome seeing or local seeing, have been studied at major observatories for decades.

In 1991, researchers at the Canada-France-Hawaii Telescope identified contributions from mirror seeing, driven by temperature differences between the primary mirror and the air immediately above it, and dome seeing, driven by temperature differences between the air inside the dome and the outside ambient air (Racine et al., 1991).

Subsequent work in 1994 at the Anglo-Australian Telescope found mirror seeing to be a more significant contributor to local seeing than dome seeing. (Ryan & Wood, 1995) This result was confirmed in 2023 by Munro et al. (2023) with optical turbulence sensing.

In today's large telescopes, thermal issues are compounded by massive structural elements which cannot equilibrate to rapidly changing ambient temperatures on the timescale of a single exposure. We set out to characterize how these persistent gradients within structural components couple with wavefront errors.

1.3. This Work

This system is complex. In general, there are rapid changes in the wavefront at short timescales while the system is spatially fixed. Since gravitational loads should be constant when the system is stationary, these are likely due to either (1) thermal effects, or (2) atmospheric seeing.

Predictable relationships between focus and dynamic thermal effects beyond the already accounted for bulk temperature contributions could be used to further refine the Look-Up Tables and improve open-loop active optics corrections. Isolating and characterizing thermal effects on seeing at the timescale of around 30 minutes (i.e. that of a single commissioning imaging experiment) would allow for the distinction between telescope-induced image degradation from atmospheric seeing and from the cosmological signals that Rubin is designed to measure.

The central exploratory research question for this work was: which thermal variables (if any) predict residual Z_4 patterns after LUT corrections on short timescales.

We assess dynamic thermal effects on Z_4 in 23 open-loop runs where elevation and azimuth are fixed. We specifically explore independent correlations and Ridge regression coefficient distributions for temperature sensors, gradient proxies, and PCA-derived thermal modes.

The remainder of this paper is structured as follows. Section 2 describes the open-loop stability run dataset, the temperature sensor suite, and the Z_4 measurement procedure. Section 3 presents the analysis methodology, including mean-centering pre-processing convention, the construction of temperature gradient proxies, the PCA decomposition, the run smoothness metric, the Savitzky-Golay slope estimation procedure, and the statistical framework for correlation testing and multiple comparisons correction. Section 4 presents the results in four stages: individual sensor correlation at the levels scale, multivariate Ridge regression, PCA-based analysis, and slope-regime correlations, followed by a comparison of dominant predictors across timescales and a stratification on results by run smoothness. Section 5 interprets these findings in the context of dome and mirror seeing theory, discusses the limitations of the linear modeling framework, and identifies direction for future analysis. Section 6 summarizes the main conclusions.

2. DATA

2.1. Dataset

The data used in this study comprises observations obtained from November 2, 2025 to December 24, 2025. The observations analyzed here are open-loop stability runs. Throughout a stability run, the telescope is held at fixed elevation, azimuth, and rotation angle. The AOS is only operated in open-loop mode, such that no real-time wavefront correction is applied in between exposures after the LUT corrections. The measured Z_4 reflects the residual defocus optical error after LUT correction.

We began with a total of 25 runs over 17 unique nights. The majority of runs consist of 40 sequences with two outliers at 19 and 44 sequences. These two runs were removed from analysis. Each run spanned approximately 28 minutes. See Table 1 for more details.

2.2. Temperature Data

Temperature data were primarily retrieved from the Engineering Facility Database (EFD), which records telemetry from sensors distributed across the telescope structure and enclosure environment. The sensors used in this analysis span five physical groupings: enclosure and ambient air, mirror glass, telescope structure, optical alignment components, and sonic anemometers. The sensor names, locations, and what they measure are detailed in Table 2.

Temperature sensor readings were queried from the EFD using the `lsst_efd_client` Python interface (Rubin Observatory/LSST, 2026). For each exposure sequence, the observation start and end times were taken from the Z_4 measurement record. A sensor reading was associated with each sequence by querying the relevant EFD time series topic over the interval $[t_{\mathrm{start}},\, t_{\mathrm{end}} + \delta]$, where $\delta = 0.2\,s$ was added to the end time to account for timestamp jitter at the boundary of the exposure window. The mean of all sensor samples falling within this window was taken as the representative temperature value for that sequence. Formally, for sequence (d, s) with observation window $[t_{\mathrm{start}}, t_{\mathrm{end}}]$, the temperature assigned to sensor A is:

$$T_{A,d,s} = \frac{1}{N_{A,d,s}} \sum_{t_k \in [t_{\mathrm{start}},, t_{\mathrm{end}}+\delta]} T_A(t_k) \qquad (5)$$

where $T_A(t_k)$ are the individual EFD samples within the window and $N_{A,d,s}$ is the number of samples available. Sequences for which no samples were returned for a given sensor were assigned `NaN` and excluded from analyses involving that sensor.

2.3. Z_4 Measurement

As mentioned in Section 1.1, the Zernike polynomials are determined through curvature wavefront sensing. This is a mathematical reconstruction of the telescope's wavefront error created through analysis of defocused star images in corner wavefront sensors.

The corner wavefront sensors on LSSTCam are split such that an intra-focal (1.5 mm in front of focus) and extra-focal (1.5 mm behind focus) image are captured simultaneously. The differences in intensity distribution between the two planes allows an algorithm to estimate the Zernike polynomial decomposition (Xin et al., 2015).

The Z_4 Zernike coefficient characterizes defocus and is used throughout this analysis as a proxy for open-loop image quality. A value of $Z_4 = 0$ corresponds to the nominal focus predicted by the LUT; deviations indicate residual defocus not accounted for by the open-loop correction. The typical within-run standard deviation of Z_4 across the dataset was $\sim 0.1\,\mu$m, while the between-run standard deviation was $\sim 0.37\,\mu$m. Z_4 varied by approximately $\sim 0.4\,\mu$m peak-to-peak within a typical

Table 1: Summary of the open-loop stability run dataset. Two runs were excluded from analysis due to data quality issues (Run 9: mid-run gap of 377 s; Run 13: truncated at 19 sequences). All statistics are reported for the 23-run analysis dataset unless otherwise noted.

Property	**Value**
Total runs collected	25
Runs excluded	2 (Runs 9, 13)
Analysis runs	23
Observation dates	11/01/2025 – 12/23/2025
Sequences per run (typical)	40
Run duration (typical)	∼28 min
Total sequences (analysis)	920
Z_4 statistics (23-run analysis dataset)	
Mean of run means ($\bar{Z}_4$ across nights)	-0.0284 μm
Night-to-night std	0.3720 μm
Overall Z_4 range (all runs)	-1.1132 to 0.8740 μm
Mean within-run σ_{Z_4}	0.0998 μm
Between-run σ_{Z_4}	0.3720 μm
Median within-run peak-to-peak	0.4030 μm

Table 2: Temperature sensors used in the correlation analysis, with physical location and display name used in figures.

Sensor	**Physical location**
Enclosure and ambient air	
Near-M2 air	Enclosure air near secondary mirror ($z \approx 6.8$ m)
Weather station	Ambient air, Young 32400 weather station outside dome
Above-M1M3 air	Air above primary mirror cell
M1M3-level air	Air at M1M3 mirror height within enclosure
Camera-level air	Enclosure air near camera ($z \approx 3.5$ m)
Mirror glass	
M1M3 glass	Primary mirror (M1M3) glass temperature
M2 glass ring 5	Secondary mirror glass, ring sensor 5
M2 glass ring 6	Secondary mirror glass, ring sensor 6
Telescope structure	
TMA truss $+x/+y$	Telescope mount assembly truss, $+x/+y$ quadrant
TMA truss $-x/-y$	Telescope mount assembly truss, $-x/-y$ quadrant
Optical alignment components	
Camera hexapod	Camera hexapod rotator structure
M2 hexapod	Secondary mirror hexapod structure
Camera rotator	Camera rotator structure
Other	
Sonic temperature	Temperature derived from sonic anemometer speed of sound

Table 3: Temperature gradient proxies used in the correlation analysis, defined as pairwise differences between sensor readings and mean-centered within runs. Gradients are grouped by physical mechanism.

Gradient	**Physical interpretation**
Enclosure-to-ambient	
Near-M2 air − Weather station	Classic dome seeing predictor. Warm dome interior drives buoyancy-forced outflow through the slit.
M1M3-level air − Near-M2 air	Vertical stratification within enclosure between mirror-level and near-M2 air.
Structural-to-enclosure-air	
TMA $-x/-y$ − Near-M2 air	Truss thermal lag vs enclosure air. Dominant slope-scale predictor.
TMA $+x/+y$ − Near-M2 air	As above for second truss sensor.
TMA $-x/-y$ − Weather station	Total structural lag relative to nighttime ambient environment.
TMA $+x/+y$ − TMA $-x/-y$	Asymmetric truss loading; can produce non-rotationally-symmetric aberrations.
TMA $-x/-y$ − M1M3-level air	Within-enclosure vertical gradient combining structural lag with mirror environment. Dominant levels-scale predictor.
Mirror glass	
M1M3 glass − M1M3-level air	Most direct mirror seeing signal. Drives surface-layer convection above the primary mirror.
M1M3 glass − Near-M2 air	Primary mirror thermal lag relative to bulk dome environment.
M2 glass ring 5 − Near-M2 air	Secondary mirror thermal lag; M2 equilibrates faster than M1M3.
M2 ring 5 − M2 ring 6	M2 glass spatial asymmetry; non-uniform M2 temperature can produce astigmatic aberrations.
Optical alignment components	
Camera hex − Near-M2 air	Camera hexapod thermal lag; retained as reference for interpreting the hexapod differential.
M2 hex − Camera hex	Differential expansion between M2 and camera hexapods directly produces a net defocus error.
Vertical enclosure gradient	
Near-M2 air − Camera-level air	Vertical thermal stratification within dome between $z \approx 6.8\,\mathrm{m}$ and $z \approx 3.5\,\mathrm{m}$.

run. Figure 3 shows six representative runs of raw Z_4 as a function of elapsed time in the run.

2.4. Data Reduction

Z_4 was averaged across wavefront sensors to produce a single `z4_mean` value per sequence. Temperature readings were similarly averaged within the exposure window for sensors with sub-exposure cadence. The resulting matched dataset was stored as a flat per-sequence table with one row per exposure, one column per sensor, and `run_id` identifying the observing sequence.

3. METHODS

This section describes the analysis pipeline applied to the open-loop stability run dataset. We begin by defining the pre-processing convention used to isolate within-run thermal and optical variation from the run-level offsets that dominate the raw data. We then explain the errors in the two excluded runs and the metric for evaluating smoothness of each run used to later stratify results. We next describe the construction of temperature gradient proxies, the principle component decomposition of the thermal sensor suite, and the Savitzky–Golay slope estimation. Finally we describe the statistical framework for correlation testing, including the autocorrelation correction and multiple comparisons adjustment.

3.1. Pre-processing

Mean-centering within runs All analysis operates on within-run residuals after subtracting the within-run mean from each sensor. This technique attempted to remove the bulk temperature affects and isolate dynamic effects of particular temperatures on focus.

Mean-centering alone did not completely remove collinearity as is reflected by the PCA analysis and the moderate to high off-diagonal correlation between sensors in Figure 4. Some components seem particularly spatially coupled such as the structural sensors of the TMA trusses and M2 hexapod. Others are surprisingly independent, such as M2 glass ring 5 and ring 6 which have a pair correlation value of only 0.07. The camera rotator sensor is negatively correlated with all other temperature sensors.

Excluded runs Two runs were excluded from analysis. Run 9, taken on November 8, 2025, had a 377-second gap in the middle of exposure data. Run 13, taken on November 11, 2025, had only 19 sequences compared to the standard 40 sequences. It was likely truncated for an unknown operational reason.

Evaluating smoothness To investigate whether thermal–optical coupling is more detectable in runs where Z_4 evolves smoothly, we classified each run as smooth or dynamic based on the lag-1 autocorrelation of the raw Z_4 sequence. For each run r, the lag-1 autocorrelation is:

$$\rho_1^{(r)} = \mathrm{Corr}\left(Z_{4,i,r}, Z_{4,i+1,r}\right) \tag{6}$$

where the correlation is taken over all adjacent pairs within the run. A high value of $\rho_1^{(r)}$ indicates that consecutive Z_4 measurements are strongly predictable from one another, reflecting a run in which Z_4 evolves gradually and coherently — either drifting monotonically or varying on timescales long relative to the 43 s cadence. A low value indicates rapid, unpredictable fluctuations in which each measurement carries little information about the next. Runs were partitioned into two groups using a threshold of $r_1^{(r)} = 0.5$, selected by inspection of the distribution of lag-1 auto-correlations across the 23 analysis runs. Runs with lag-1 autocorrelation larger than 0.5 were classified as smooth; those with lag-1 autocorrelation less than 0.5 were classified as dynamics. This yielded 17 smooth runs and 6 dynamic runs.

3.2. Thermal Gradient Proxies

Raw temperature sensor readings are not the most physically motivated predictors of optical aberration. Dome seeing and thermally-induced structural deformation are driven by temperature *differences* between components — the temperature differential between the telescope structure and the enclosure air, or between the mirror surface environment and the ambient — rather than by the absolute thermal level of any single sensor. We therefore construct a set of pairwise temperature differential proxies as synthetic predictor variables, defined as the difference between two sensor readings at the same timestamp:

$$\Delta T_{AB,i} = T_{A,i} - T_{B,i} \tag{7}$$

where $T_{A,i}$ and $T_{B,i}$ are the readings of sensors A and B at observation i. Each gradient is mean-centered within runs using the same convention applied to raw sensors. Fourteen gradients were defined *a priori* based on physical reasoning about the dominant dome seeing mechanisms identified in the literature and the sensor groupings revealed by the PCA decomposition. The selection criterion was that each gradient should capture a physically interpretable differential between two functionally distinct telescope regions: enclosure air versus ambient, structural components versus enclosure air, mirror glass versus surrounding air, or hexapod assemblies versus enclosure air. Gradients were organized into three physical groups reflecting these mechanisms: enclosure-to-ambient gradients, which capture the classical dome seeing mode; structural-to-enclosure-air gradients, which capture the thermal lag of the telescope structure relative to the dome environment; and mirror and hexapod gradients, which capture differential thermal effects

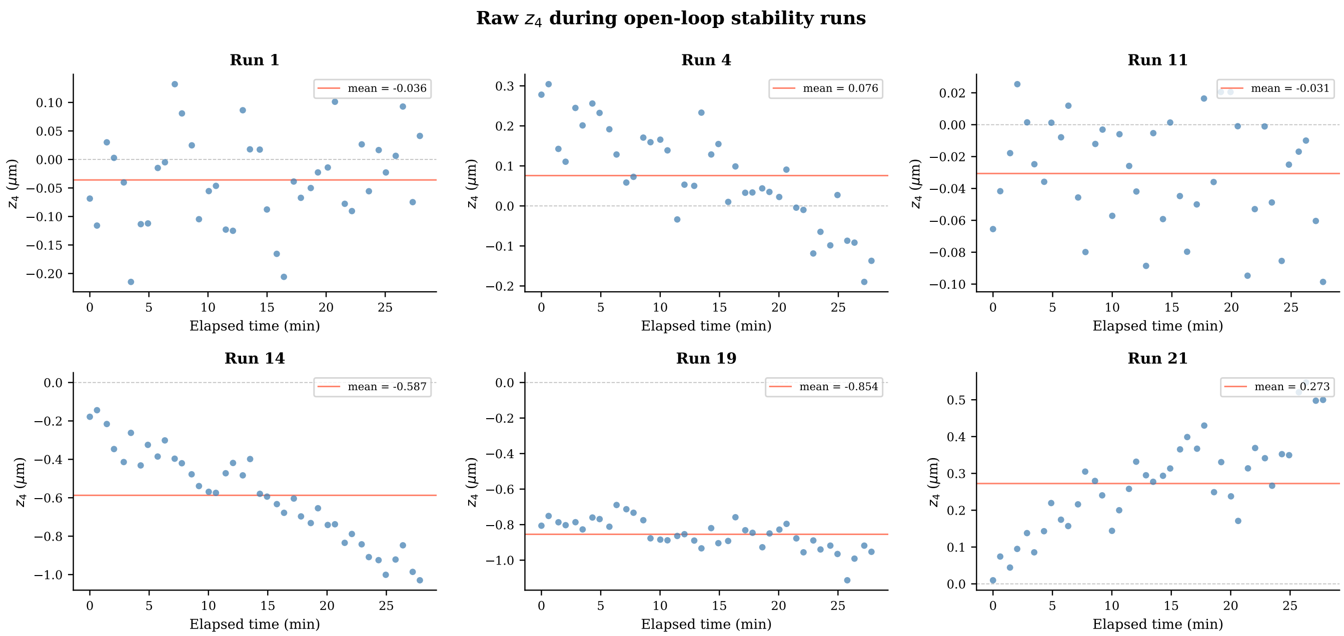

Figure 3: Raw Z_4 (defocus Zernike coefficient) as a function of elapsed time during six representative open-loop stability runs. Each point represents a single exposure sequence. The red horizontal line marks the within-run sample mean of Z_4. Runs 1 and 11 are classified as dynamic under later smoothness evaluation, while runs 4, 14, 19, and 21 are classified as smooth.

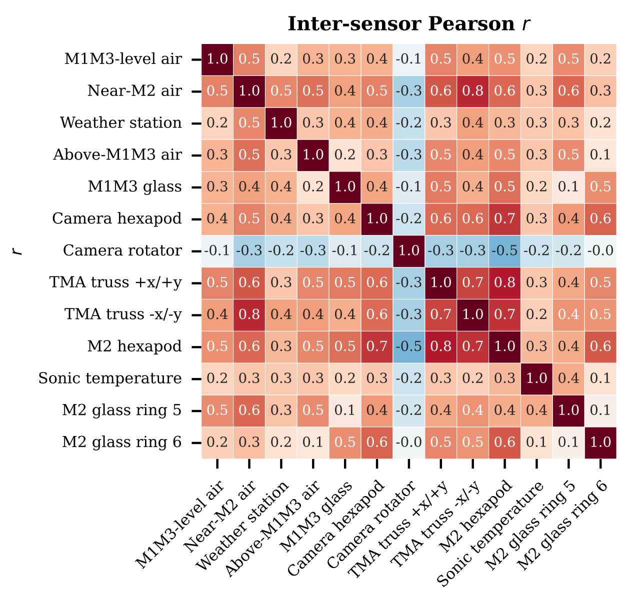

Figure 4: Pairwise Pearson correlations between all mean-centered temperature sensors across the 23 analysis runs. High off-diagonal values indicate collinearity between sensors.

in the optical components most directly responsible for defocus. Full definitions, sensor pairs, and physical descriptions for all fourteen gradients are given in Table 3.

3.3. Single Sensor Correlation Metrics

Quenouille lag-1 effective sample size Adjacent observations within a run are generally auto-correlated, both because temperature changes slowly relative to the sampling cadence and because the Savitzky & Golay (1964) differentiation procedure induces artificial auto-correlation by sharing input points across adjacent slope estimates. Using the raw sample size n in significance tests would therefore overstate statistical power. We apply the Quenouille (1949) lag-1 effective sample size correction, which estimates the number of approximately independent observations as

$$N_{eff} = \frac{n(1 - r_1^{(x)} r_1^{(y)})}{(1 + r_1^{(x)} r_1^{(y)})} \tag{8}$$

where $r_1^{(x)}$ and $r_1^{(y)}$ are the lag-1 auto-correlations of the predictor and target series respectively. All reported p-values use N_{eff} in place of n.

Benjamini-Hochberg (BH) correction Because each analysis involves simultaneous correlation tests across multiple predictors, uncorrected p-values would overstate the evidence due to the multiple comparisons problem. We apply the Benjamini-Hochberg (BH) procedure (Benjamini & Hochberg, 1995) to control the false discovery rate (FDR) at $\alpha = 0.05$.

3.4. Regression

To assess the independent contribution of each temperature sensor and gradient proxy to $Z_4^{(c)}$, we fit multivariate regression models relating the full set of predictors to the optical target simultaneously. This complements the pairwise correlation analysis by controlling for inter-predictor collinearity, at the cost of requiring the assumption that relationships between thermal sensors and Z_4 is linear.

OLS Ordinary least squares (OLS) regression minimizes the sum of squared residuals:

$$\hat{\boldsymbol{\beta}}_{\mathrm{OLS}} = \arg\min_{\boldsymbol{\beta}} |\mathbf{y} - \mathbf{X}\boldsymbol{\beta}|_2^2 \tag{9}$$

where $\mathbf{y} \in \mathbb{R}^n$ is the vector of mean-centered Z_4 observations and $\mathbf{X} \in \mathbb{R}^{n\times p}$ is the standardized predictor matrix. All predictors were standardized to zero mean and unit variance within each run prior to fitting, so that coefficients are directly comparable across sensors with different temperature ranges. This method does not control for collinearity, thus is only used on regression over the principle components from PCA because they are already orthogonal.

Ridge regression To stabilize coefficient estimates in the presence of collinearity, we fit Ridge regression, which augments the OLS objective with an l_2 penalty on the coefficient vector:

$$\hat{\boldsymbol{\beta}}_{\mathrm{Ridge}} = \arg\min_{\boldsymbol{\beta}} |\mathbf{y} - \mathbf{X}\boldsymbol{\beta}|_2^2 + \alpha\,|\boldsymbol{\beta}|_2^2 \tag{10}$$

The penalty term $\alpha||\boldsymbol{\beta}||_2^2$ shrinks all coefficients toward zero. In the presence of collinear predictors, Ridge distributes predictive credit across correlated variables rather than allocating it arbitrarily to whichever happens to be numerically dominant in a given sample. The regularization parameter was selected independently for each run by cross-validation over $\alpha \in \{0.01, 0.1, 1.0, 10.0\}$.

3.5. Principle Component Analysis

Principal component analysis (PCA) was applied to the thermal sensor suite to characterize the dominant modes of temperature variation within the dataset and to construct orthogonal predictors of Z_4 that are free from the inter-sensor collinearity that limits individual sensor regression. PCA was performed on the standardized sensor matrix $\mathbf{X} \in \mathbb{R}^{n\times p}$, where each column was standardized to zero mean and unit variance using a `StandardScaler` fit across all observations. The eigenvectors of the sample covariance matrix of $\mathbf{X}$ define the principal component (PC) loading vectors $\mathbf{v}_1, \ldots, \mathbf{v}_p$, ordered by decreasing explained variance. The score of observation i on component k is the projection $s_{k,i} = \mathbf{x}_i^\top \mathbf{v}_k$. The PCA decomposition was fit on mean-centered sensor readings $\mathbf{x}_i^{(c)}$ and captures only the within-run differential thermal structure that is directly comparable to the mean-centered $Z_4^{(c)}$ target. PC scores were additionally mean-centered within runs after projection, ensuring that the variance alignment convention applied throughout the analysis is maintained. The first K = 9 principal components were retained for analysis, accounting for 93.6% of total within-run thermal variance.

3.6. Slope Estimation

To investigate whether thermal dynamics predict optical dynamics independently of the quasi-static thermal state, we estimate the instantaneous first derivative of each time series within each run.

Cadence correction The observation cadence alternates systematically between approximately 36s and 50s across all runs. Applying Savitzky-Golay differentiation directly to this non-uniform cadence would introduce a systematic derivative bias of approximately 17-23% at every other point, since the filter assumes uniform spacing. We therefore resample all time series onto a uniform 43s grid via linear interpolation prior to differentiation. Linear interpolation was chosen because the interpolation step is close to the original cadence, minimizing approximation error, and because missing values propagate cleanly as `NaN` rather than being filled by extrapolation.

Savitzky-Golay differentiation The first derivative of each resampled time series is estimated using a Savitzky-Golay (SG) filter with window length $w = 7$ points (corresponding to $7 \times 43\,\mathrm{s} \approx 5\mathrm{min}$) and polynomial order $p = 3$. The SG filter jointly smooths and differentiates by fitting a local polynomial to a sliding window, avoiding the amplification of high-frequency noise inherent in finite differencing. The filter was applied independently within each run and the $\lfloor w/2 \rfloor = 3$ edge points at each end of each run were set to `NaN`.

Parameter sensitivity Primary results were replicated with window lengths $w \in \{5, 7, 9\}$. Qualitative conclusions were unchanged across this range; $w = 7$ is reported throughout as it provides sufficient smoothing to suppress single-exposure noise while resolving thermal dynamics on timescales of a few minutes.

4. RESULTS

The analysis is presented in six sections which build on the statistical methods and findings of the prior analysis. First, we examine the relationship between each temperature sensor and Z_4 at the *levels scale* – exploring whether the thermal state of the telescope (absolute level of a sensor reading, mean-centered within each run) predicts the optical state ($Z_4^{(c)}$). We extend this analysis to multivariate regression, with methods discussed in Section 3.4, to assess whether sensors carry independent information beyond their pairwise correlations. Next, we directly test the hypothesis that particular proxies of thermal gradients reflecting dome and mirror seeing impact Z_4 with the proxies detailed in Section 3.2. Next, we use PCA to decompose the thermal state into orthogonal modes and examine how they couple with Z_4. Then, these analyses are repeated at the *slope scale* – exploring whether the rate of change of temperature predicts the rate of change of $Z_4^{(c)}$. Finally, we attempt to stratify and compare the analysis by the smoothness of the run, characterized by methods discussed in Section 4.6.

All correlations are reported as Pearson r with autocorrelation-corrected p-values (Quenouille lag-1 N_{eff} correction) and Benjamini-Hochberg multiple comparisons adjustment at $\alpha = 0.05$, both discussed in Section 3.3. Significance throughout refers to survival of the BH correction unless otherwise stated.

4.1. Single Sensor Correlations

We begin by computing the Pearson correlation between each mean-centered temperature sensor and mean-centered Z_4 across all 23 analysis runs pooled. This establishes which individual sensors carry linear predictive information about the optical state of the quasi-static (levels) timescale when we consider all runs together.

Figure 5(a) shows the pooled linear correlation strength r between each temperature sensor and $Z_4^{(c)}$. The only temperature sensors which survive the BH significance correction are

- M2 Glass Ring 6 ($r = 0.292$),
- TMA Truss -x/-y ($r = 0.255$),
- M1M3 Glass ($r = -0.207$), and
- Camera Hexapod ($r = 0.204$)

listed in descending absolute r value. Scatter plots of each of the top four sensors against $Z_4^{(c)}$ colored by run are presented in Figure A1. These show that the linear best fit is not representative of all runs, with some in the opposite direction of the pooled correlation. A detailed table of the correlation metrics is presented in Table 4.

Table 4: Contemporaneous correlation metrics between mean-centered temperature sensors and Z_4. All p-values are autocorrelation-corrected using the Quenouille effective sample size. † indicates Pearson-Spearman discrepancy > 0.03, suggesting outlier sensitivity. Bold indicates significance under Benjamini-Hochberg correction.

Sensor	r	ρ	$\lvert r-\rho \rvert$	p
M2 glass ring 6	**0.292**	**0.251**	0.041†	**0.0000**
TMA truss -x/-y	**0.255**	**0.232**	0.023	**0.0002**
M1M3 glass	**-0.207**	**-0.032**	0.175†	**0.0024**
Camera hexapod	**0.204**	**0.148**	0.056†	**0.0028**
M2 hexapod	0.135	0.062	0.073†	0.0400
Weather station	-0.110	-0.074	0.036†	0.0954
Near-M2 air	0.112	0.168	0.056†	0.0999
M1M3-level air	-0.109	-0.043	0.066†	0.1173
TMA truss +x/+y	0.072	0.076	0.004	0.2922
Sonic temperature	-0.052	-0.087	0.035†	0.2960
M2 glass ring 5	0.063	0.036	0.027	0.3172
Camera rotator	-0.033	0.021	0.054†	0.6263
Above-M1M3 air	-0.004	-0.059	0.055†	0.9527

We can also compare the Pearson r value with the Spearman ρ, as presented in Table 4, to assess whether the correlations are driven by outliers. Close agreement between the two metrics would indicate that the trend is monotonic. Disagreements of larger than around 0.05 flag outlier sensitivity. Among our four BH significant sensors, only M2 glass ring 6 and TMA truss -x/-y have an $|r - \rho|$ value less than 0.05.

The pairwise correlations alone cannot determine whether a sensor that correlates with $Z_4^{(c)}$ has independent predictive information from the general collinearity of the sensors. A sensor may appear to couple with Z_4 when in fact both Z_4 and the sensor are coupled with a separate overall thermal mode driving the effect. To address this, we turn to multivariate Ridge regression.

4.2. Multivariate Regression

We now attempt multivariate Ridge regression, which estimates each sensor's contribution to $Z_4^{(c)}$ while holding all other sensors consistent. Ridge regression also regularizes coefficients, shrinking redundancy towards zero to handle collinearity.

We run Ridge regression on each run independently. The distribution of coefficients per temperature across each of the 23 analysis runs is presented in Figure 6. Consistency in the distribution on one side of zero would indicate a reliable relationship between a given sensor and Z_4. All of the individual thermal sensor Ridge regression coefficient distributions cross zero, indicating no consistent relationship between individual sensors and $Z_4^{(c)}$. There is fairly wide spread in the distributions of each sensors' coefficients across all runs and many outliers.

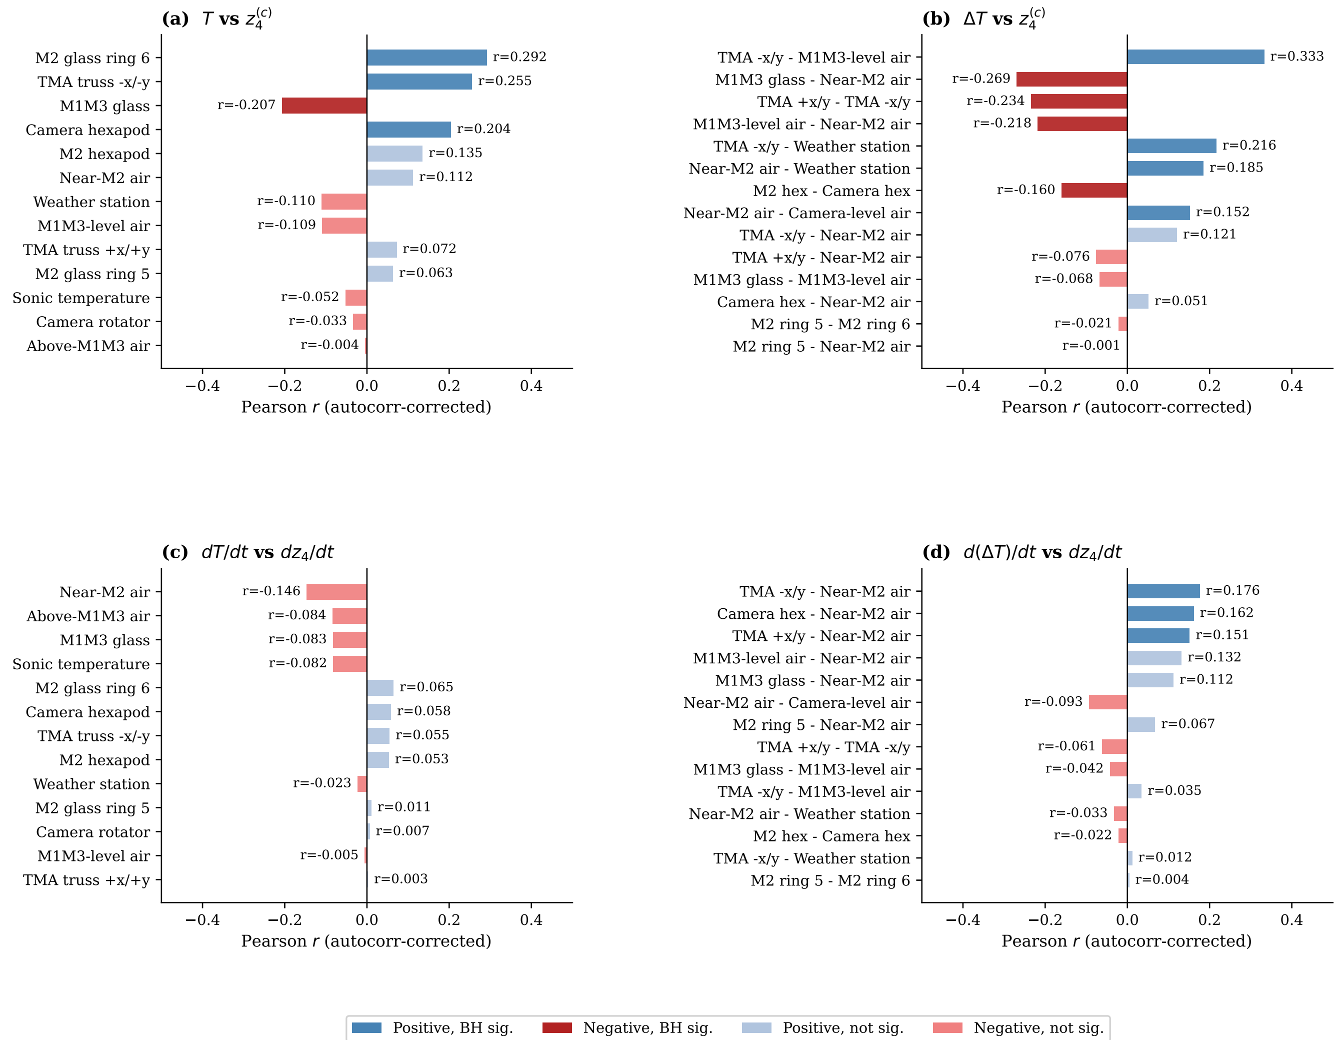

Figure 5: Pearson r correlations between thermal predictors and Z_4 across all 23 analysis runs, with autocorrelation-corrected p-values (Quenouille lag-1 $N_{\rm eff}$ correction) and Benjamini–Hochberg multiple comparisons adjustment at $\alpha = 0.05$. Solid bars survive BH correction; faded bars do not. Blue bars indicate positive correlation, red bars negative. (a) Individual temperature sensor levels vs mean-centered $Z_4^{(c)}$. (b) Temperature gradient proxy levels vs $Z_4^{(c)}$. (c) Individual sensor slopes (dT/dt) vs dZ_4/dt, estimated via Savitzky-Golay differentiation on a uniform 43 s grid. (d) Gradient proxy slopes ($d(\Delta T)/dt$) vs dZ_4/dt.

4.3. Gradient Proxy Correlations

We then repeated the individual correlation analysis looking at thermal gradients as proxies for dome and mirror seeing. A table of these 14 gradients, which sensors they involve, and their physical meanings are presented in Table 3.

Figure 5(b) shows the correlation strength of each gradient proxy with $Z_4^{(c)}$. Eight gradients survive BH significance correction. These are

- TMA Truss -x/-y – M1M3-level Air ($r = 0.333$),
- M1M3 Glass – Near-M2 Air ($r = -0.269$),
- TMA Truss +x/+y – TMA Truss -x/-y ($r = -0.234$),
- M1M3-level Air – Near-M2 Air ($r = -0.218$),
- TMA Truss -x/-y – Weather Station ($r = 0.216$),
- Near-M2 Air – Camera Hexapod ($r = -0.160$), and
- Near-M2 Air – Camera-level Air ($r = 0.152$)

listed in descending order of $|r|$. Table 5 presents the correlation metrics of the gradient proxies and $Z_4^{(c)}$. All measures except for M1M3-level Air – Near-M2 Air are flagged for discrepancy between Pearson r correlation

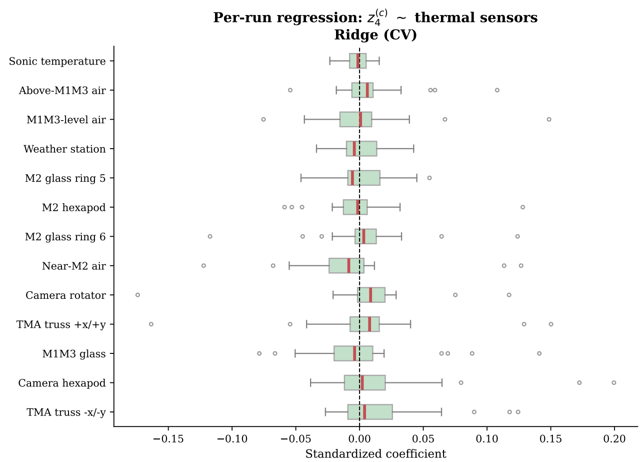

Figure 6: Distribution of per-run standardized regression coefficients from Ridge cross-validated regression of $Z_4^{(c)}$ on all mean-centered temperature sensors simultaneously. Each box summarizes the coefficient for one sensor across 23 runs. The red median line indicates the central tendency.

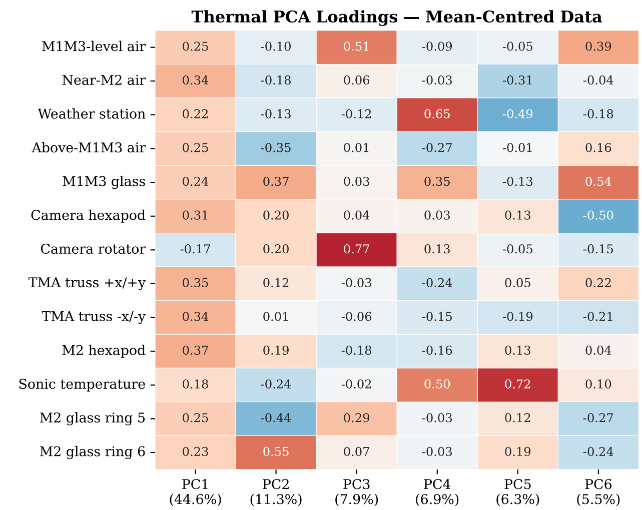

Figure 7: PCA loading heatmap for the first five principle components computed on mean-centered sensor data.

metric and Spearman ρ, indicating that most of the correlation between gradients and $Z_4^{(c)}$ is sensitive to outliers.

Table 5: Contemporaneous correlation metrics between mean-centered temperature gradient proxies and Z_4. All p-values are autocorrelation-corrected using the Quenouille effective sample size. † indicates Pearson-Spearman discrepancy > 0.03, suggesting outlier sensitivity. Bold indicates significance under Benjamini-Hochberg correction.

Gradient	r	ρ	$\lvert r-\rho \rvert$	p
TMA $-x/y$ − M1M3-level air	**0.333**	**0.269**	0.064†	**0.0000**
M1M3 glass − Near-M2 air	**-0.269**	**-0.204**	0.065†	**0.0001**
TMA $+x/y$ − TMA $-x/y$	**-0.234**	**-0.120**	0.114†	**0.0008**
TMA $-x/y$ − Weather station	**0.216**	**0.173**	0.043†	**0.0008**
M1M3-level air − Near-M2 air	**-0.218**	**-0.192**	0.026	**0.0013**
Near-M2 air − Weather station	**0.185**	**0.153**	0.032†	**0.0041**
M2 hex − Camera hex	**-0.160**	**-0.240**	0.079†	**0.0212**
Near-M2 air − Camera-level air	**0.152**	**0.233**	0.081†	**0.0258**
TMA $-x/y$ − Near-M2 air	0.121	0.096	0.024	0.0691
TMA $+x/y$ − Near-M2 air	-0.076	-0.014	0.062†	0.2683
M1M3 glass − M1M3-level air	-0.068	-0.021	0.047†	0.3259
Camera hex − Near-M2 air	0.051	0.015	0.037†	0.4590
M2 ring 5 − M2 ring 6	-0.021	-0.098	0.077†	0.7459
M2 ring 5 − Near-M2 air	-0.001	-0.074	0.073†	0.9858

4.4. Principle Component Correlation

A recurrent confounding consideration in this analysis is the collinearity of thermal sensors. As the global thermal mode changes or particular spatially coupled sensors change together, it is difficult to isolate individual effects on seeing. PCA decomposition of the thermal mode separates orthogonal components which maximally explain thermal variance. This is discussed in further detail in Section 3.5.

Figure 7 presents the PCA loadings – coefficients of the linear combination of variables which compose each principle component – computed on the mean-centered thermal sensor data. Principle Component 1 (PC1) captures almost half (44.6%) of the thermal variance in the mean-centered data and is composed of a fairly equal positive contribution of all thermal sensors, other than Camera Rotator which has a negative contribution. This seems to be residual variance due to the overall thermal mode which is not completely removed by mean-centering.

Table 6 gives the correlation strength between each mean-centered principal component score and $Z_4^{(c)}$. The correlations of PC 4, 6, and 8 individually with $Z_4^{(c)}$ survive BH correction. They have r values of -0.261, -0.427, and 0.197 respectively. PC 4 is flagged for Pearson-Spearman discrepancy, indicating sensitivity to outliers.

Scatter plots of each of the first six principle components against $Z_4^{(c)}$ along with the correlation score r, autocorrelation-correction p-value, and line of best fit are given in Figure A2.

Controlling for collinearity by using orthogonal principle components, we can use OLS regression to regress $Z_4^{(c)}$ on a linear combination of the PCs. Figure 8 represents the standardized OLS coefficients for each principle component on data pooled across all 23 analysis runs. The R^2 value, which represents the portion of variance in $Z_4^{(c)}$ which is explained by the linear combination of PCs, is fairly low ($R^2 = 0.33$).

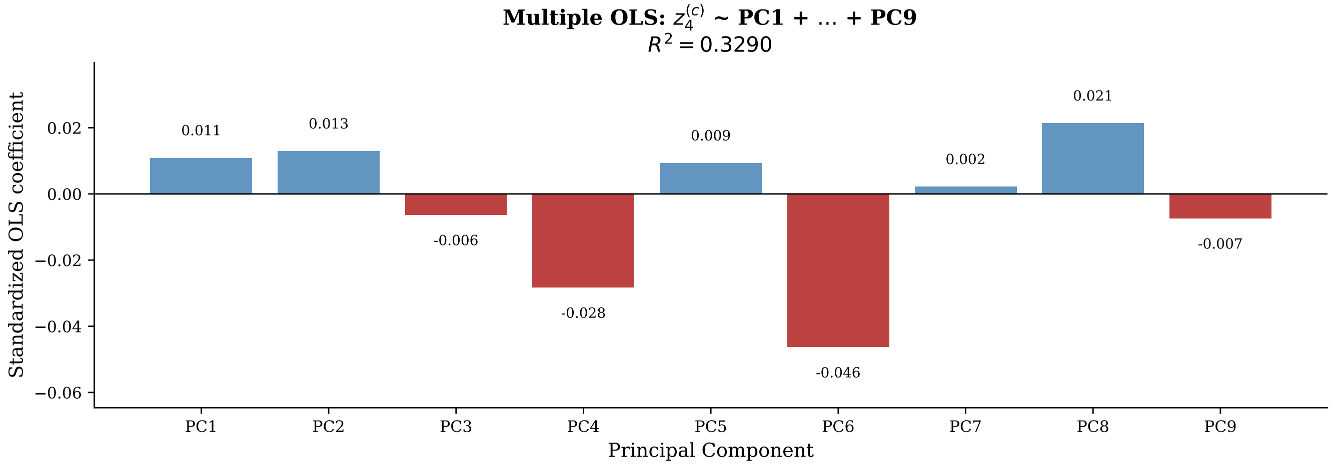

Figure 8: Multiple OLS regression of $Z_4^{(c)}$ on all five mean-centered PC scores simultaneously. Bars show standardized coefficients; the reported R^2 indicates the fraction of within-run Z_4 variance explained jointly by the thermal PC modes.

Table 6: Pearson r and Spearman ρ between each mean-centered principal component score and mean-centered Z_4, with autocorrelation-corrected p-values (Quenouille lag-1 N_{eff} correction) and Benjamini-Hochberg multiple comparisons adjustment. PC scores are derived from the mean-centered thermal PCA basis and additionally mean-centered within runs prior to correlation. † indicates Pearson-Spearman discrepancy > 0.03, suggesting outlier sensitivity. Bold indicates significance under Benjamini–Hochberg correction.

Principle Component	Variance %	Pearson r	Spearman ρ	p value
PC 1	44.6	0.1103	0.1276	0.1380
PC 2†	11.3	0.1191	0.1803	0.0828
PC 3†	7.9	-0.0579	0.0007	0.4041
PC 4†	6.9	**-0.2606**	**-0.1902**	**0.0000**
PC 5	6.3	0.0858	0.0657	0.1158
PC 6	5.5	**-0.4270**	**-0.4007**	**0.0000**
PC 7	4.5	0.0198	0.0364	0.7624
PC 8	3.5	**0.1971**	**0.1707**	**0.0021**
PC 9	3.1	-0.0675	-0.0506	0.2917

4.5. Slope Analysis

The mean-centered levels analysis embeds a core tension. The practice of mean-centering supposes prior knowledge about the run itself, particularly how any immediate measure compares to the future aggregate. We have also seen in our principle component analysis that mean-centering does not completely remove global thermal structure. Ideally, we would isolate the impact of a given temperature change to the change in Z_4 independent of the global thermal mode. To do so, we compare the slope of Z_4 with respect to time at each point with the change in thermal sensor readings at each point. The mathematical calculation of the slope and additional computational details are given in Section 3.6.

Temperature Sensors Figure 5(c) shows the correlation strength of the slope of each temperature sensor (dT/dt) with the slope of Z_4 (dZ_4/dt). None of the direct thermal sensor slopes had BH significant correlations with dZ_4/dt or surpassed an absolute r value of 0.15. Table 7 presents the Pearson r correlation strength and p-values for each temperature sensor in the levels timescale and slope timescale for comparison.

Table 7: Pearson r and autocorrelation-corrected p-values for individual temperature sensors correlated with $Z_4^{(c)}$ (levels) and dZ_4/dt (slope). Bolded values denote predictors surviving Benjamini–Hochberg correction at $\alpha = 0.05$. Sorted by $|r|$ at levels scale.

	$Z_4^{(c)} \sim T^{(c)}$		$\dot{z}_4 \sim \dot{T}$	
Sensor	r	p	r	p
M2 glass ring 6	**+0.292**	**0.000**	+0.065	0.269
TMA truss -x/-y	**+0.255**	**0.000**	+0.055	0.342
M1M3 glass	**-0.207**	**0.002**	-0.083	0.171
Camera hexapod	**+0.204**	**0.003**	+0.058	0.335
M2 hexapod	**+0.135**	**0.040**	+0.053	0.359
Near-M2 air	+0.112	0.100	-0.146	0.006
Weather station	-0.110	0.095	-0.023	0.667
M1M3-level air	-0.109	0.117	-0.005	0.928
TMA truss +x/+y	+0.072	0.292	+0.003	0.966
M2 glass ring 5	+0.063	0.317	+0.011	0.831
Sonic temperature	-0.052	0.296	-0.082	0.080
Camera rotator	-0.033	0.626	+0.007	0.909
Above-M1M3 air	-0.004	0.953	-0.084	0.108

Gradient Proxies Figure 5(d) shows the correlation strength of the slope of each temperature gradient sensor ($d(\Delta T)/dt$) with the slope of Z_4 (dZ_4/dt). Three of the gradient proxies have BH-significant correlations with Z_4 at the slope scale. These are

- TMA Truss -x/-y – Near-M2 Air ($r = 0.176$),
- Camera Hexapod – Near-M2 Air ($r = 0.162$), &
- TMA Truss +x/+y – Near-M2 Air ($r = 0.151$).

Table 8: Pearson r and autocorrelation-corrected p-values for temperature gradient proxies correlated with $Z_4^{(c)}$ (levels) and $\dot{z}_4$ (slope). Bold values indicate predictors surviving Benjamini–Hochberg correction at $\alpha = 0.05$. Sorted by $|r|$ at levels scale.

	$Z_4^{(c)} \sim \Delta T^{(c)}$		$\dot{z}_4 \sim d(\Delta T)/dt$	
Gradient	r	p	r	p
Truss -x/-y − M1M3-level air	**+0.333**	**0.000**	+0.035	0.548
M1M3 glass − Near-M2 air	**-0.269**	**0.000**	+0.112	0.040
Truss +x/+y − Truss -x/-y	**-0.234**	**0.001**	-0.061	0.279
M1M3-level air − Near-M2 air	**-0.218**	**0.001**	+0.132	0.016
Truss -x/-y − Weather station	**+0.216**	**0.001**	+0.012	0.826
Near-M2 air − Weather station	**+0.185**	**0.004**	-0.033	0.537
M2 hexapod − Camera hexapod	**-0.160**	**0.021**	-0.022	0.714
Near-M2 air − Camera-level air	**+0.152**	**0.026**	-0.093	0.086
Truss -x/-y − Near-M2 air	+0.121	0.069	**+0.176**	**0.001**
Truss +x/+y − Near-M2 air	-0.076	0.268	**+0.151**	**0.005**
M1M3 glass − M1M3-level air	-0.068	0.326	-0.042	0.471
Camera hexapod − Near-M2 air	+0.051	0.459	**+0.162**	**0.003**
M2 glass ring 5 − M2 glass ring 6	-0.021	0.746	+0.004	0.932
M2 glass ring 5 − Near-M2 air	-0.001	0.986	+0.067	0.181

Table 8 compares the predictive strength of each temperature gradient proxy at the levels and slope timescales. The gradients are sorted in descending order of $|r|$ at the levels scale. The eight BH-significant gradients with the largest levels scale correlation are not significant at the slope scale. However, the next two largest levels correlations are significant at the slope scale. No gradient is significant at both scales.

4.6. Stratify Runs on Smoothness

Another potential confounding pattern is the consistency of Z_4 throughout the run. Notice in Figure 3, runs 4, 14, 19, and 21 seem to have relativity smooth scatter and a consistent trend over time while runs 1 and 11 are very noisy with no consistent trend. Perhaps this noise is interfering with and masking thermal coupling. To test this hypothesis, we will stratify the runs by smooth or dynamic using the metric described in Section 3.1.

Figure 9 plots the per-run Pearson $|r|$ between the thermal sensor and $Z_4^{(c)}$ as a function of the smoothness of each run (lag-1 autocorrelation of Z_4). Plots are given for each of the top five sensors in the single sensor pooled correlation analysis. A positive trend would support the hypothesis that smooth runs are signal-dominated while dynamic runs are noise-dominated and obscure thermal-optical coupled. This seems to be consistently true in Figure 9. We see a statistically significant positive relationships between the smoothness measure and Pearson r correlation metric for all five of the top correlation sensors.

Figure 10 shows the distribution of per-run standardized Ridge regression coefficients stratified by the smoothness of the run. For each of the thermal variables, the smooth runs have wider variance in coefficient distribution. The dynamic runs are more narrow in variance, but still cluster near zero without a striking consistent predictive value for any of the temperature sensors. There are still no consistent trends in sign or magnitude for any of the thermal variables in the smooth or dynamic groups of runs.

5. DISCUSSION

The central finding of this analysis is that the thermal–optical system at the Rubin Observatory resists simple characterization, even in the controlled setting of open-loop commissioning runs. No single temperature sensor or gradient proxy has more than modest correlation with Z_4 or maintains a consistent relationship with Z_4 throughout the runs analyzed. This analysis reveals some structured patterns which suggest the dynamic thermal–optical relationship is distributed across many interacting mechanisms and timescales. Linear models and independent sensors correlation do not robustly capture the system's complexity nor account for the collinearity of sensors.

5.1. Independent Predictors

Of the 13 temperature sensors, 10 have Pearson-Spearman discrepancies suggesting sensitivity to outliers. Only one of the four sensors with significant correlation to Z_4 has sufficiently low Pearson-Spearman discrepancy to seem resistant to outlier noise. This was TMA Truss -x/-y. Despite this seeming significant relationship, plotting TMA Truss -x/-y against $Z_4^{(c)}$ in Figure A1 shows that many of the runs have relationships between TMA Truss -x/-y and $Z_4^{(c)}$ not well fit by the correlation value.

Given the general disagreement between Pearson r and Spearman ρ and the stark difference in linear trends depending on the run, individual sensors do not seem to robustly and dependably predict pooled Z_4 trends over the course of all runs.

This finding is to be expected. Given the complexity of the system and the interdependence of sensors, one temperature independently and consistently explaining dynamic Z_4 behavior would be surprising. This is best shown with the variability of Ridge regression coeffi-

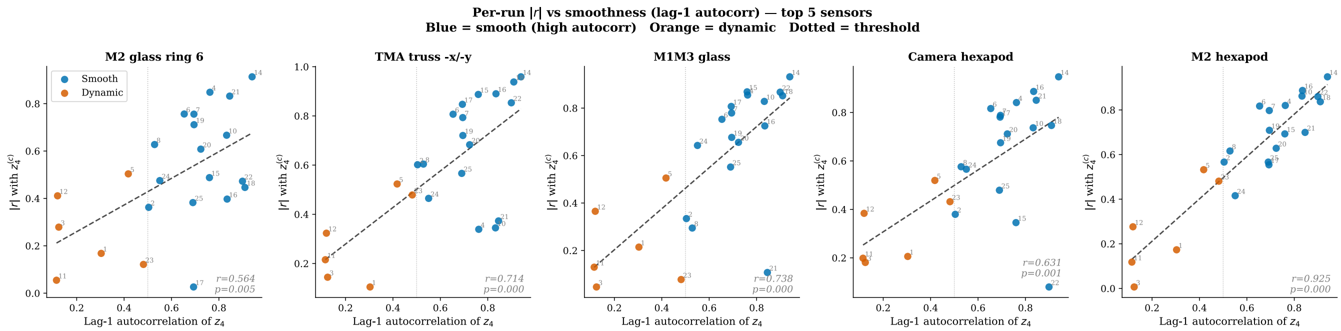

Figure 9: Per-run Pearson $|r|$ between the top thermal sensor and mean-centered Z_4 plotted against the run smoothness metric.

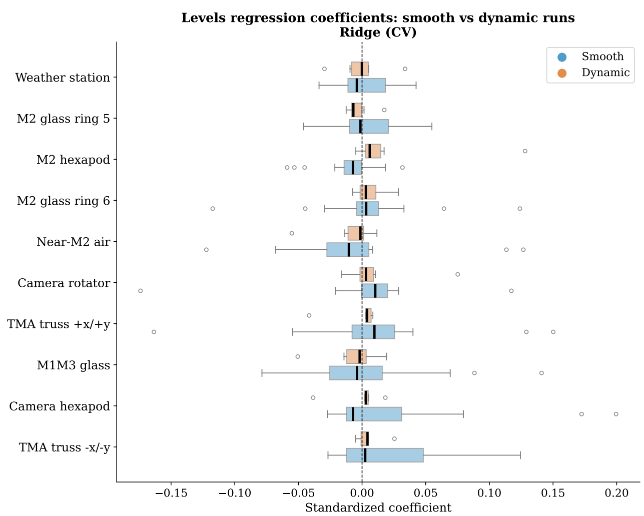

Figure 10: Distribution of per-run standardized regression coefficients for the levels regression $Z_4^{(c)}$ on all mean-centered temperature sensors, stratified by smoothness group. OLS results are show on the left, Ridge cross-validated results on the right.

cients. The inter-quartile range of coefficient distribution crosses zero for every sensors. This indicates the lack of a consistent independent predictor variable. Though, for some of our significantly correlated sensors, the distribution is skewed. Camera Hexapod and TMA Truss -x/-y both have strong positive coefficient outliers for a few of the runs. This shows that in some particular runs these variables are highly predictive, but those are particular non-representative cases.

The same patterns are broadly true for the temperature differential proxy measurements as well. None are strong independent predictors. A larger number of gradients have BH-corrected significant correlations to $Z_4^{(c)}$, but again all but one have Pearson-Spearman discrepancy indicating outlier sensitivity. Interestingly, this single significant gradient whose Pearson-Spearman discrepancy is sufficiently small is M1M3 Air – Near-M2 Air. This has a stronger correlation than traditional dome seeing between inside and outside air. The gradient throughout the air column is more spatially close.

The strongest correlation is between inside air near M1M3 and one of the truss structures. The gradient between the other truss structure sensor and M1M3 air was not included in this analysis, but was checked to see if this relationship is generalizable. It had a statistically significant but weaker correlation at $r = 0.181$. Perhaps that discrepancy is reasonable given the predictive strength of the gradient between the two truss structures of $r = -0.234$.

The second strongest correlation at $r = -0.269$ seems to reflect more traditional mirror seeing on first glance, however it relates the M1M3 mirror with the air near M2. Compare this with the weakest of our correlations, the M2 mirror and the air near M2. This puzzling details again undermine the interpretability of these correlations and suggest much of it is likely due to general noise and outlier sensitivity.

5.2. PCA

The collinearity of sensors might be handled by first decomposing them into orthogonal thermal modes through PCA. The first principle component, PC1, explains 44.6% of the mean-centered thermal sensor variance and is comprised of a relatively equal positive contribution from all temperature sensors (excluding Camera Rotator, which has a negative contribution to PC1). This indicates that some of the overall thermal mode is not completely removed by mean-centering alone. Future work could use detrending by subtracting the first principle component from the raw data to better isolate dynamic thermal modes. Further principle components have interesting compositions with point to other gradients and sensor interactions.

The second principle component has prominent positive contributions from M1M3 glass and M2 glass ring 6 and prominent negative contributions Above-M1M3 air and M2 glass ring 5 with moderate contributions from many of the other variables. This seems to combine two gradients – one between the two sensors on the M2 glass and another between the M1M3 glass and air above it. PC2 explains 11.3% of the remaining variance in the thermal suite.

PC3 has a very strong positive contributions from Camera Rotator and M1M3-level air with loadings of 0.77 and 0.51 respectively.

PC4 seems to capture outside temperature with contributions of 0.65 from the weather station sensor and 0.50 from the sonic temperature. There are also some moderate negative contributions from air above M1M3 and structural sensors. This may capture dome seeing effects between the ambient air temperature outside the dome and the inside air coupled with structural component temperatures. PC4 has a moderate anti-correlation with $Z_4^{(c)}$ at BH-corrected significant levels, but is flagged for Pearson-Spearman discrepancy.

PC5 seems to be primarily capturing the difference between the weather station sensor and sonic temperature, but is not well correlated with $Z_4^{(c)}$.

PC6 has the strongest Pearson correlation with $Z_4^{(c)}$ of $r = -0.427$ and is not flagged for outlier sensitivity. The strongest positive loadings of PC6 are M1M3 glass and air with a prominent negative loading from the camera hexapod. This may indicate a thermal gradient throughout the height of the telescope in both the air column and the structure itself.

Though the principle components point to interesting combinations of effects, they are still not robustly linearly predictive of $Z_4^{(c)}$. Using OLS regression gives a sub-par R^2 value of 0.329.

5.3. Slope and Time Scales

The correlation analysis between slopes of temperature sensors and Z_4 measures implicitly controls for overall thermal state, looking specifically about how changes in temperature state impact the change of Z_4. There were no statistically significant correlations between the slope of temperature sensors over time with immediate change in Z_4.

Near-M2 air was included in all of the top 5 strongest gradient correlations at the slope time scale. None of these correlations were particularly strong with the largest value being only 0.176 for TMA Truss -x/-y – Near-M2 Air. No gradient proxy was significant at both timescales.

5.4. Noise

Stratification between runs which are particularly noisy and those which follow a more consistent trend might help isolate between unexplained or external noise and predictable thermally driven variation. Plotting our smoothness measure against correlation strength with $Z_4^{(c)}$ shows a statistically significant positive trend across all five strongest predictors. This indicates that smoother runs are more directly coupled with thermal effects. The dynamic runs likely have other sources of noise not accounted for in this analysis causing the more volatile movement of Z_4.

However, the distribution of smooth run ridge regression coefficients do not show consistent robust independent predictors for any of the temperature sensors. The smooth runs generally have wider distributions than the dynamics ones. This was a somewhat surprising trend. Although the absolute correlation metric may be stronger in smooth runs, the direction of the strength is not consistent.

This supports the idea that Z_4 and independent thermal variables may be moving independently in most runs, but couple by chance in a few, creating the artifact of strong prediction in variable direction between runs.

5.5. A Core Tension

Many of the statistical and data science methods used in this analysis are not directly interpretable. The coefficients in Ridge regression give an idea of the relevant variables and their effect on predictions, but they do not tell a definitive cause and effect story. The principle components explain variance and seem to relate to our focus measure, but they are not meaningfully interpretable either. We can optimize for accuracy in making predictions of focus, but these effects are difficult to isolate and finding the true thermal impactors rather than proxies or noise through the project of prediction and correlation is a complex task.

5.6. Limitations

The largest limitation in isolating thermal effects is the collinearity between sensors. It is practically impossible isolate thermal change in an individual component to validate that it drives focus aberrations.

Although Ridge regression handles collinearity, presenting the distribution in this fashion may mask some genuine trends. It is possible that in one run, the Ridge regression arbitrarily allocates predictive preference to one collinear sensor over the others. The opposite could be true in another run. This would cause the distribution to cluster around zero. Sensors which demonstrate this pattern of high individual correlation with $Z_4^{(c)}$ but unstable Ridge coefficients could be interpreted as belonging to a collinear group that collectively predicts $Z_4^{(c)}$.

The slope approximation method used in this analysis lowers the effective sample number significantly. Other methods may allow for accurate optimization while maintaining more data points.

5.7. Future Work

Future analysis should explore more complex non-linear modeling. This could be in the form of extended nonlinear statistical methods like the ones explored in this analysis, or modern machine learning, though these require much more data than is currently available and are not easily interpretable.

The slope correlation analysis would benefit from considering lagging impacts. The rate of change is a very immediate measure. These effects might be better understood with small delays between the change in temperature and the potential correlated change in Z_4. This analysis considers small windows of time around each point in fitting a polynomial and calculating the slope, but further analysis might explore explicit lag analysis.

The same analysis pipeline repeated with different temperature sensors could probe other effects we are currently blind to. Extending to other sensors in the telescope structure could refine our understanding structural thermal gradient effects. The camera skin temperature might be of particular interest

This analysis pipeline could also be replicated with other seeing targets, such as higher order Zernike coefficients.

6. CONCLUSIONS

The Rubin Observatory system is complex. Dynamic thermal effects on focus are not adequately captured by independent variable regression and linear modeling. The statistical methods in this analysis are inherently limited by lack of interpretability and isolation of particular physical effects. Future work could include additional temperature sensors, gradient proxies, and seeing targets, as well as benefit from more complex non-linear modeling.

REFERENCES

Benjamini, Y., & Hochberg, Y. 1995, Journal of the Royal Statistical Society: Series B, 57, 289

Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, The Astrophysical Journal, 873, 111, doi: 10.3847/1538-4357/ab042c

Munro, J., Hansen, J., & Travouillon, T. 2023, Journal of Astronomical Telescopes, Instruments, and Systems, 9, 017004, doi: 10.1117/1.JATIS.9.1.017004

Noll, R. J. 1976, Journal of the Optical Society of America, 66, 207, doi: 10.1364/JOSA.66.000207

Poyneer, L. A., Troy, M., Macintosh, B., & Gavel, D. T. 2003, Optics Letters, 28, 798, doi: 10.1364/OL.28.000798

Quenouille, M. H. 1949, Journal of the Royal Statistical Society, Series B, 11, 68

Racine, R., Salmon, D., Cowley, D., & Sovka, J. 1991, Publications of the Astronomical Society of the Pacific, 103, 1020, doi: 10.1086/132920

Rubin Observatory/LSST. 2026, EFD Client: A Python Interface to the Engineering and Facility Database, https://efd-client.lsst.io/index.html

Ryan, S. G., & Wood, P. R. 1995, Publications of the Astronomical Society of Australia, 12, 89

Savitzky, A., & Golay, M. J. E. 1964, Analytical Chemistry, 36, 1627, doi: 10.1021/ac60214a047

Welford, W. T. 1986, Aberrations of Optical Systems (Bristol: Adam Hilger)

Xin, B., Claver, C., Liang, M., et al. 2015, Applied Optics, 54, 9045, doi: 10.1364/AO.54.009045

A. SUPPLEMENTARY FIGURES

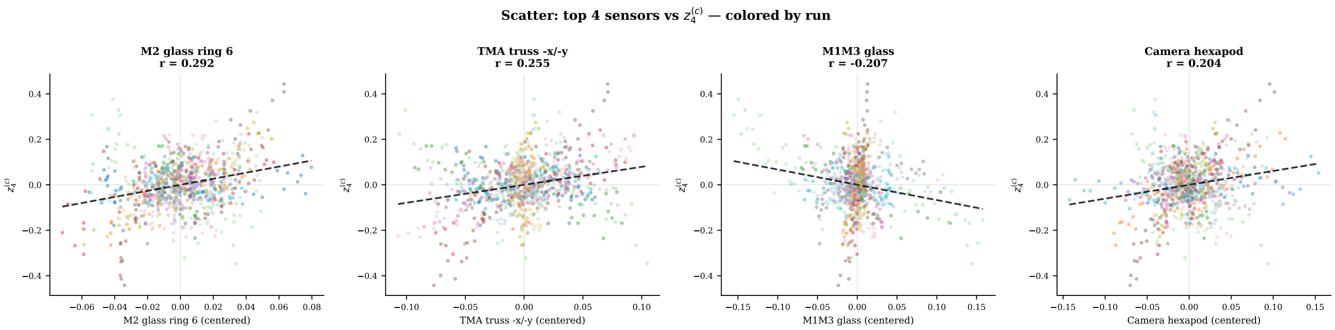

Figure A1: Scatter plots of the top four sensors against $Z_4^{(c)}$ at levels scale, colored by run. The dashed line shows the pooled OLS fit.

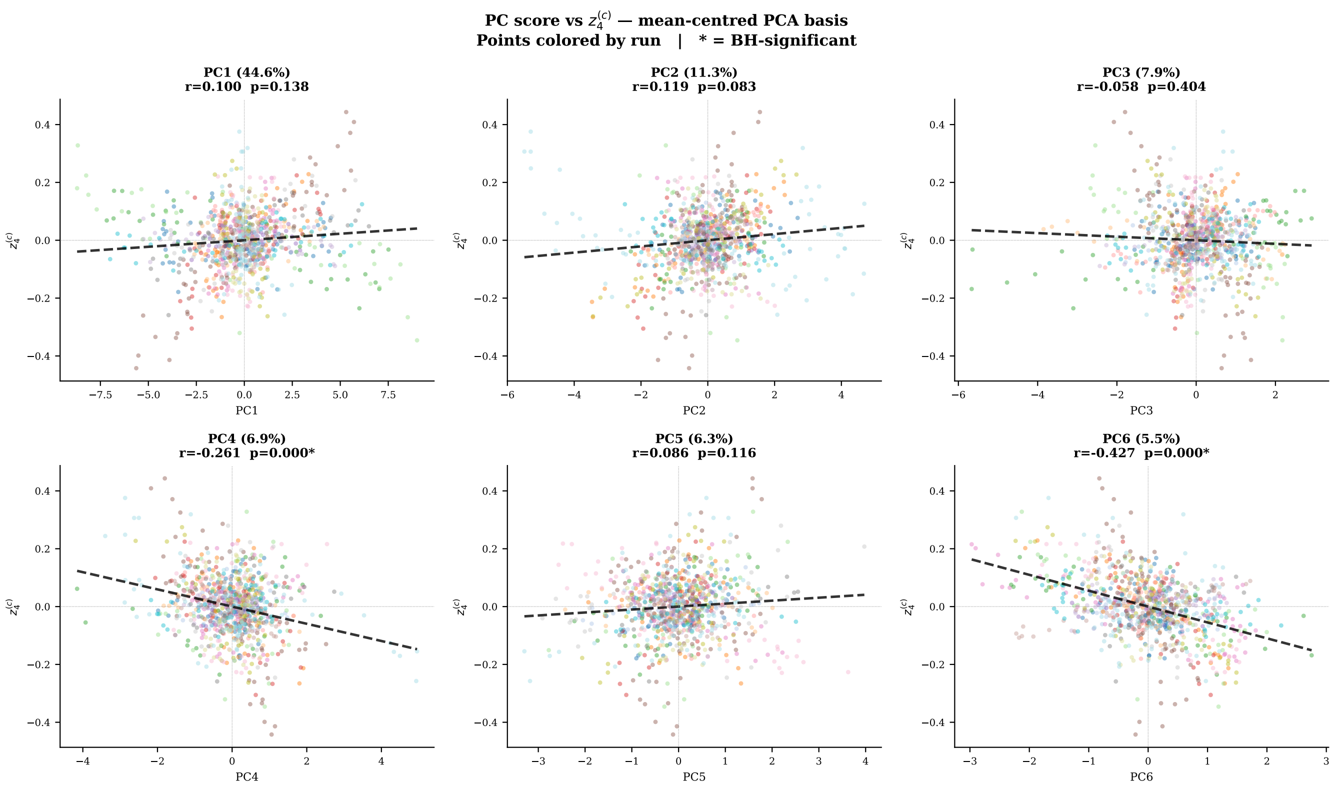

Figure A2: Scatter plots of mean-centered PC scores against $Z_4^{(c)}$. Points are colored by run. The dashed line shows the pooled OLS fit. Pearson r, effective sample size N_{eff}, and autocorrelation-corrected p-value are reported in each panel title. Asterisks mark components surviving BH correction.

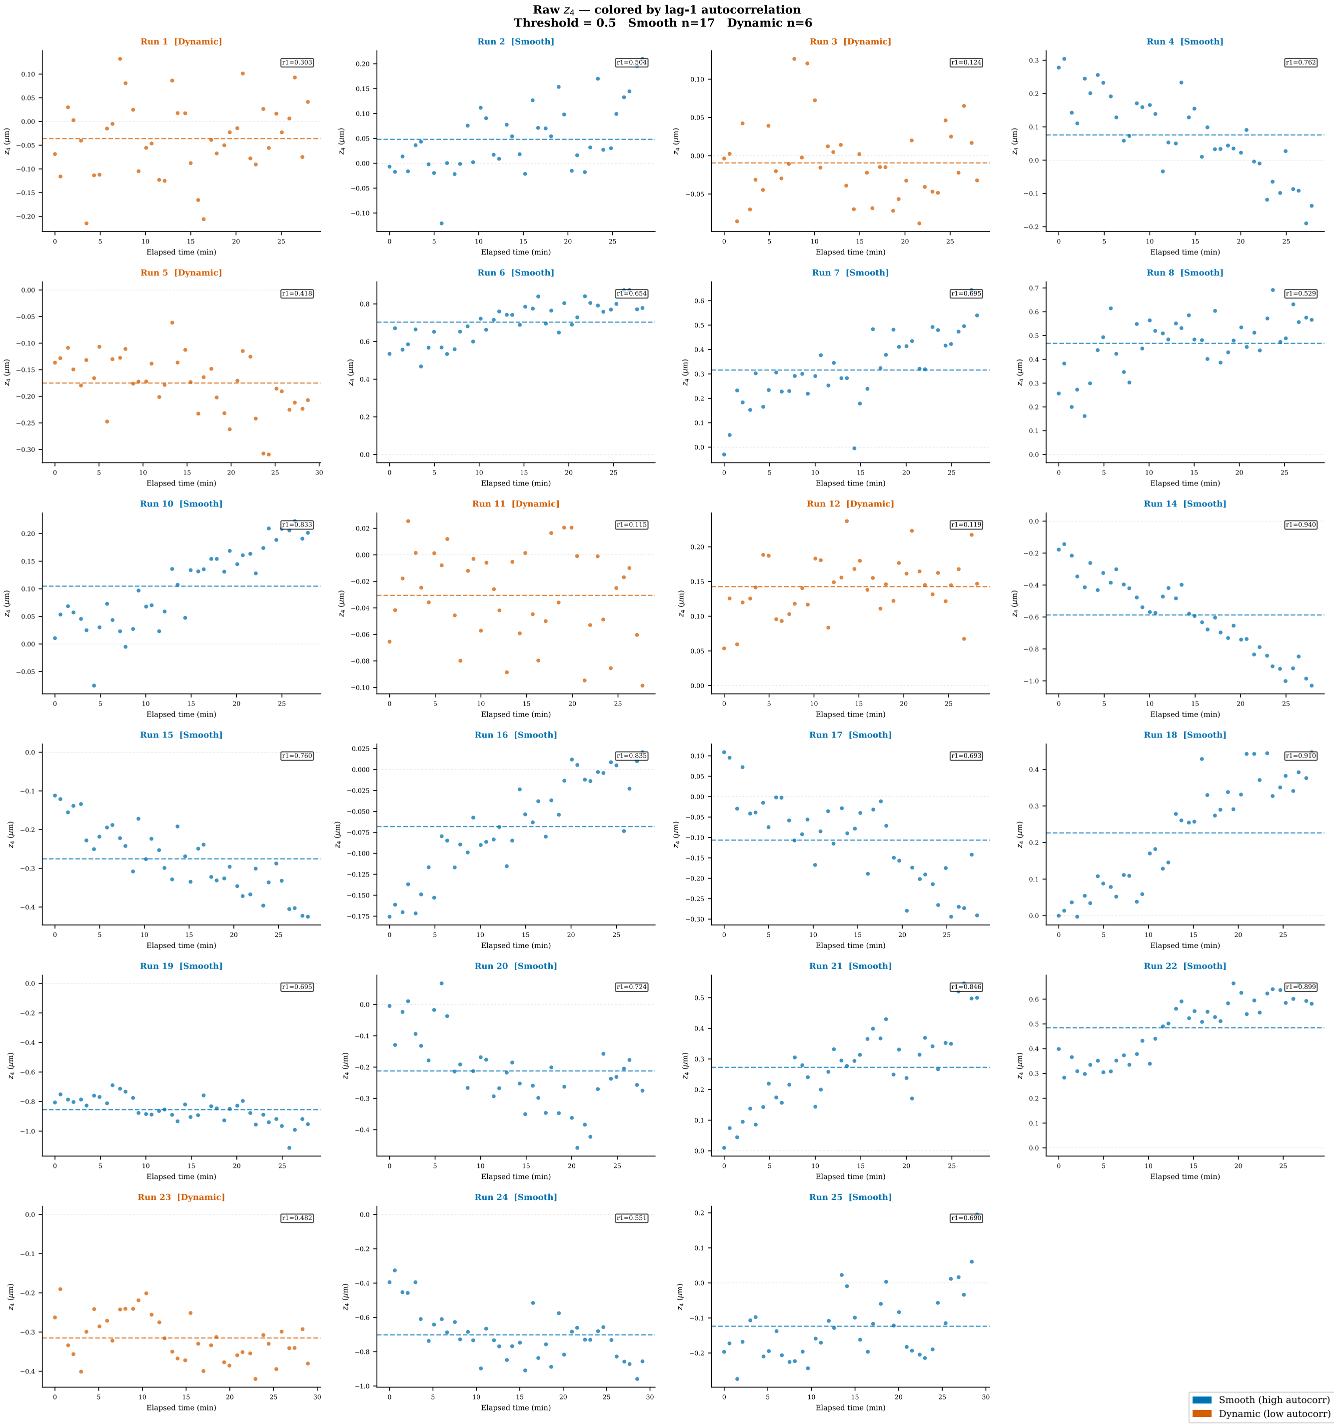

Figure A3: Raw Z_4 as a function of elapsed time for all 23 analysis runs, colored by smoothness classification. Blue runs have high lag-1 autocorrelation ($\rho_1^{(r)} \geqslant 0.5$), indicating that consecutive Z_4 measurements are strongly predictable from one another and that Z_4 evolves coherently within the run. Orange runs have low lag-1 autocorrelation ($\rho_1^{(r)} < 0.5$), indicating rapid or irregular Z_4 fluctuations. The dashed horizontal line marks the within-run mean. The lag-1 autocorrelation value $\rho_1^{(r)}$ is reported in the upper right corner of each panel.

Table A1: Summary of open-loop stability runs used in this analysis. Runs 9 and 13 were excluded from all analyses (indicated by †). Day obs is the observation date in MM/DD/YYYY format. $\bar{z}_4$ denotes the within-run mean defocus coefficient, σ_{Z_4} the within-run standard deviation, and ΔZ_4 the peak-to-peak range.

Run	**Day obs**	$N_{\mathbf{seq}}$	**Duration (min)**	$\bar{z}_4$ (μ**m**)	σ_{Z_4} (μ**m**)	ΔZ_4 (μ**m**)
1	11/01/2025	40	27.9	-0.0359	0.0804	0.3468
2	11/01/2025	40	28.4	0.0480	0.0696	0.3307
3	11/02/2025	40	28.0	-0.0093	0.0500	0.2149
4	11/03/2025	40	27.8	0.0758	0.1220	0.4941
5	11/04/2025	40	28.7	-0.1750	0.0552	0.2480
6	11/05/2025	40	28.2	0.7031	0.1017	0.4062
7	11/05/2025	40	27.8	0.3160	0.1452	0.6739
8	11/07/2025	40	28.0	0.4670	0.1190	0.5299
9†	11/07/2025	44	36.7	0.3744	0.1421	0.4784
10	11/09/2025	40	27.9	0.1048	0.0722	0.2981
11	11/09/2025	40	27.7	-0.0306	0.0345	0.1241
12	11/10/2025	40	28.2	0.1427	0.0418	0.1834
13†	11/10/2025	19	12.9	0.1296	0.0645	0.2817
14	11/10/2025	40	27.8	-0.5873	0.2409	0.8855
15	11/11/2025	40	28.3	-0.2755	0.0862	0.3131
16	11/11/2025	40	27.9	-0.0680	0.0580	0.1963
17	11/12/2025	40	28.4	-0.1066	0.1056	0.4030
18	11/13/2024	40	28.2	0.2263	0.1488	0.4497
19	11/14/2025	40	27.8	-0.8544	0.0850	0.4237
20	11/15/2025	40	27.8	-0.2122	0.1199	0.5254
21	11/17/2025	40	27.8	0.2725	0.1303	0.5371
22	11/21/2025	40	28.2	0.4848	0.1212	0.3804
23	12/11/2025	40	28.9	-0.3151	0.0591	0.2291
24	12/23/2025	40	29.1	-0.7015	0.1505	0.6330
25	12/23/2025	40	29.0	-0.1237	0.0979	0.4693
All 25 runs						
23 analysis runs						