Draft version September 12, 2026
Typeset using LaTeX **modern** style in AASTeX7.0.1

TURBO-GP: Removing atmospheric turbulence from the Vera C. Rubin Observatory Astrometry through scalable Gaussian Processes.

Pierre-François Léget [1]

[1] *Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA*

(Dated: September 12, 2026)

ABSTRACT

Description of the application of TUrbulence Removal By fourier-Optimized Gaussian Process on DP2 data set and compare performance to default GP used in DP2.

1. INTRODUCTION

Astrometry refers to the determination of the position of astronomical sources on the sky. For an optical wide-field telescope such as the Vera C. Rubin Observatory (Ž. Ivezić et al. 2019) and its CCD mosaic camera, LSSTCam (SLAC National Accelerator Laboratory & NSF-DOE Vera C. Rubin Observatory 2025), obtaining an astrometric solution therefore means determining the transformation that maps pixel coordinates on a given CCD to a position on the sky, known as the World Coordinate System (WCS). A precise and unbiased astrometric solution is important for a broad range of science cases, from solar system science, such as searches for trans-Neptunian objects (TNOs; P. H. Bernardinelli et al. 2020), to dark energy constraints from cosmic shear (P. F. Léget et al. 2021). The Rubin Observatory has a strict requirement on its astrometric repeatability (the dispersion of the measured positions of the same sources from one visit to another) of 10 mas, with a stretch goal of 5 mas (Ž. Ivezić & The LSST Science Collaboration 2018).

To first order, obtaining the astrometric solution consists in fitting the optical distortion that maps pixel coordinates to sky coordinates. To do so, the Rubin Observatory uses the same astrometric model as the Dark Energy Survey (DES; G. M. Bernstein et al. 2017, hereafter B17), adapted to the needs of the Rubin Observatory within the LSST Science Pipelines (Rubin Observatory Science Pipelines Developers et al. 2026), as described in C. Saunders (2024, hereafter S24), with The Monster (P. S. Ferguson et al. 2025) as the reference catalog. The astrometric residuals discussed in this paper are the difference between the position of a source fitted by this model over all the overlapping visits and its measured position in a given image. Once the optical distortion is fit, at least three effects still affect this mapping: differential chromatic refraction (DCR; see, e.g., A. A. Plazas & G. Bernstein 2012), static shifts caused by CCD defects such as tree rings (see, e.g., H. Y. Park et al. 2017), and atmospheric turbulence (see, e.g., B17). All these effects are present at the same time, but we do not discuss DCR in this analysis. We focus on the astrometric distortions induced by atmospheric turbulence, and we take a deeper look at the static patterns, which can be derived directly once atmospheric turbulence is corrected.

B17 noticed that, once the astrometric model is fit, the astrometric residuals of a single visit are compatible with a pure E-mode field with a variance of $\sim$100 mas^2. Because the residual field is pure E-mode and shows long stripe-like patterns, it was interpreted as the signature of atmospheric turbulence: the displacement field follows the gradient of the variations of the optical refractive index integrated along the line of sight. Given the nature of turbulence, all the information contained in the astrometric residual field can be described by its two-point statistics, and Gaussian Process Regression (GPR; C. E. Rasmussen & C. K. I. Williams 2006) is therefore the optimal model for it. W. F. Fortino et al. (2021, hereafter F21) modeled it with a GPR, using an isotropic von Kármán kernel and imposing a curl-free constraint on the displacement field, and were able to remove most of it, at the cost of an expensive hyperparameter optimization and inference that require multiple Cholesky decompositions. P. F. Léget et al. (2021, hereafter L21) found the same type of pattern on the Hyper Suprime-Cam (HSC): a pure E-mode field, with a variance of $\sim$30 mas^2. L21 also used a GPR with a von Kármán kernel, and demonstrated on HSC data that it performs better than the standard Gaussian kernel; but they introduced three major changes with respect to F21. First, L21 did not impose the curl-free constraint and modeled each component of the displacement field independently. Second, L21 fit the anisotropy of the spatial correlation, which varies from one exposure to another. The third change is algorithmic. Instead of estimating the hyperparameters of the kernel with a standard maximum-likelihood approach, they measured the two-point correlation function of the astrometric residual field directly and fit the hyperparameters of the von Kármán kernel to it with a simple least-squares minimization. This reduces the cost of the hyperparameter optimization from $O(N^3)$ to $O(N \log N)$, while keeping a Cholesky factorization for the inference. This method removes about 90% of the variance that turbulence introduces in the astrometric dispersion. Following the path paved by L21, D. C. H. Gomes et al. (2025, hereafter G25) pushed the approach further on DES data by skipping the hyperparameter fit altogether, as well as the constraint of an analytical function to describe the spatial correlation, and using an interpolated version of the measured two-point correlation function directly as the kernel. As there is no direct comparison with the other methods, it appears to perform similarly to the previous ones. The price is a singular value decomposition (SVD) to remove the negative eigenvalues, since the resulting covariance matrix is not positive definite. Even with this SVD, we found G25 to be faster than L21, because there is no hyperparameter fit anymore and no evaluation of the von Kármán kernel, whose calls to Bessel functions are slow. The authors claimed that the SVD is not necessary and that a Cholesky decomposition can be used instead; we were not able to reproduce their result on Rubin data without an SVD. More recently, C. Yang et al. (2026, hereafter Y26) also observed a pure E-mode pattern in astrometric residuals, but did not follow the path of F21, L21, or G25 and left the correction for further analyses.

Although these corrections improve the global astrometric solution of DES and HSC, they are not as important for those surveys as they are for the Rubin Observatory. Indeed, the amount of E-mode power introduced in a single image scales with the inverse of the square root of the exposure time, which explains why HSC shows a better overall performance than DES: the typical exposure time was 270 s for HSC, compared with 90 s for DES. The Rubin Observatory has set its exposure time to 30 s, and L21 predicted that, if what was observed on HSC is rescaled to the Rubin exposure time, a variance of 270 mas^2 should be observed. This can now be verified, as the first batch of data of the Rubin Observatory has been published as Data Preview 2 (DP2; NSF-DOE Vera C. Rubin Observatory 2026), and we will show that the observed E-mode is slightly higher than what L21 anticipated.

In this paper, we present how atmospheric turbulence actually affects the astrometric performance, leveraging the DP2 data set of the Rubin Observatory, and the strategy used to model it in the DP2 production, which is based on L21. The DP2 processing revealed the limitations of the method proposed in L21, and we decided, for future processing, to follow the road paved by G25 and to introduce significant changes that improve the performance of all aspects of the algorithm for future data releases: the measured two-point correlation function is used directly as the kernel, the linear algebra is solved in Fourier space with a covariance that is positive semi-definite by construction, the field of view no longer needs to be split into patches, and the computing time and memory footprint fit within the budget of the Rubin production. We refer to this new model as TUrbulence Removal By fourier-Optimized Gaussian Process (TURBO-GP).

The paper is structured as follows. Section 2 describes the data set and how the astrometric residual field is derived. Section 3 describes the basics of GPR, the L21 implementation used for DP2, and the new TURBO-GP algorithm. Section 4 shows the overall performance of the L21 method and of TURBO-GP, assesses how well turbulence is under control in DP2, and shows how TURBO-GP will outperform the current methodology in future releases. We then discuss how the astrometric residual field correlates with actual properties of the Earth's atmosphere, and how static CCD defects can potentially be derived directly from the GPR corrections.

2. THE DATA PREVIEW 2 AND ASTROMETRIC RESIDUALS

The data set used here is based on the Data Preview 2 (DP2; NSF-DOE Vera C. Rubin Observatory 2026); a full description of the data reduction and of what to expect from this data set can be found in Vera C. Rubin Observatory Team (2026). The key elements are the following. DP2 consists of data taken at the Rubin

Observatory with its wide-field optical camera, LSSTCam, between April 2025 and January 2026, in six optical bands (u, g, r, i, z, and y), during its commissioning and early operations phases. It contains 29 756 science-grade exposures covering about 15 000 deg^2, processed with the LSST Science Pipelines (Rubin Observatory Science Pipelines Developers et al. 2026).

The astrometric calibration of DP2 follows S24. After single-visit processing, the isolated point sources (unresolved, unblended, with a signal-to-noise ratio between 8 and 1000) shared by all the visits of a given band that overlap a given HEALPix level-3 region of the sky are associated, and a joint astrometric model is fit to them with the `gbdes` package (G. M. Bernstein 2022), using The Monster (P. S. Ferguson et al. 2025) as the reference catalog. The model is the composition of a static map per detector, from pixel coordinates to an intermediate frame, of a map per visit, from the intermediate frame to the plane tangent to the telescope boresight, and of the deterministic mapping from the tangent plane to the sky; the proper motion and parallax of the objects are fit as part of the solution. The astrometric residuals studied in this paper are, for each source of a given visit, the difference between the position of the object fitted by this model and the measured position of the source in the image, projected onto the tangent plane of the visit.

In DP2, these residuals are the input of an additional step, new with respect to DP1, that models the atmospheric turbulence of each visit with a Gaussian process following L21 (Sect. 3.2), and appends the prediction to the astrometric model as a two-dimensional spline (Vera C. Rubin Observatory Team 2026). In this paper, we take the residuals produced by the `gbdes` fit of DP2 and re-run this Gaussian process step, with the L21 model used in production and with TURBO-GP, so that both models are compared on exactly the same inputs.

This analysis is done on a subset of 28 517 visits, which are the visits with a successful atmospheric turbulence fit with the L21 method described in Sect. 3.2. Some visits overlap several HEALPix regions and therefore have several astrometric solutions. In the current processing, a visit that overlaps multiple HEALPix regions goes through as many L21 fits as there are overlapping regions, and the best astrometric solution is chosen for this visit further down the pipeline. In the same spirit, for a visit that overlaps multiple HEALPix regions, we keep only the (visit, HEALPix region) pair with the lowest E-mode variance before any atmospheric modeling, in order to avoid duplicates.

An example of what the astrometric residuals look like is shown for three visits in Figure 1. For all three visits, the astrometric residual field shows the expected stripe patterns, with spatial correlations that are curl-free (E-mode dominated). The E- and B-mode correlation functions averaged over the 28 517 visits are shown in Figure 2: the field is E-mode dominated, with an E-mode variance of $\sim$340 mas^2. This is 10 times higher than what was observed with HSC, and 3 times higher than what is observed with DES. In L21, a forecast of the expected E-mode variance as a function of exposure time was made, and the DP2 value does not fall far from the L21 prediction (270 mas^2). Figure 3 shows the L21 forecast compared with DES (F21), WFST (Y26), and the Rubin Observatory (DP2). This high E-mode variance is mostly driven by the shorter exposure time, which makes its correction more important for the Rubin Observatory than for previous surveys.

In the following section, we review the basic concepts of Gaussian Processes, describe how we implemented the L21 GP for DP2 and discuss some of its limitations, and then present the new TURBO-GP designed to address these limitations.

3. GAUSSIAN PROCESS MODELS

As discussed in Sect. 1, the displacement field induced by atmospheric turbulence can be described as a stationary Gaussian random field: its statistical properties do not depend on the position in the focal plane, only on the separation between two points, and they are entirely described by its two-point correlation function. Gaussian Process Regression is the natural tool to interpolate such a field from the positions where it is measured, the sources of a visit, to any other position of the focal plane, because it is built on exactly this description. In this section, we first recall the basic concepts of GPR (Sect. 3.1), then describe the L21 model as it was implemented for the DP2 production and the limitations we met with it (Sect. 3.2), and finally present TURBO-GP (Sect. 3.3). Throughout this paper, the two components of the astrometric residual field are modeled as two independent scalar Gaussian processes, following L21; G25 tested a joint model of the two components including their cross-correlation and found no improvement for four times the computing cost.

3.1. *Basic intro to Gaussian Process*

A Gaussian process is a probability distribution over functions (C. E. Rasmussen & C. K. I. Williams 2006). In our case the function is one component $y(\mathbf{X})$ of the astrometric residual field, where $\mathbf{X}$ is the position in the tangent plane of the visit. We assume that y is a Gaussian random field with zero mean, since the astrometric model has already been subtracted, and that it is stationary: the covariance between the values of the field at two positions depends only on their separation,

$$\xi(\mathbf{X}' - \mathbf{X}) \equiv \mathrm{Cov}\left[y(\mathbf{X}),\, y(\mathbf{X}')\right]. \tag{1}$$

The function ξ is the two-point correlation function of the field; in the GP literature it is called the kernel, and $\xi(0)$ is the variance of the field.

The field is measured at N positions $\mathbf{X}_i$, which form the training sample, with measurement uncertainties σ_i: $y_i = y(\mathbf{X}_i) + \epsilon_i$, where ϵ_i is a Gaussian noise of variance σ_i^2, independent from one source to another. The covariance matrix of these measurements is

$$C_{ij} = \xi(\mathbf{X}_i - \mathbf{X}_j) + \delta_{ij}\,\sigma_i^2, \qquad \text{i.e.} \qquad \mathbf{C} = \mathbf{K} + \mathbf{D}, \tag{2}$$

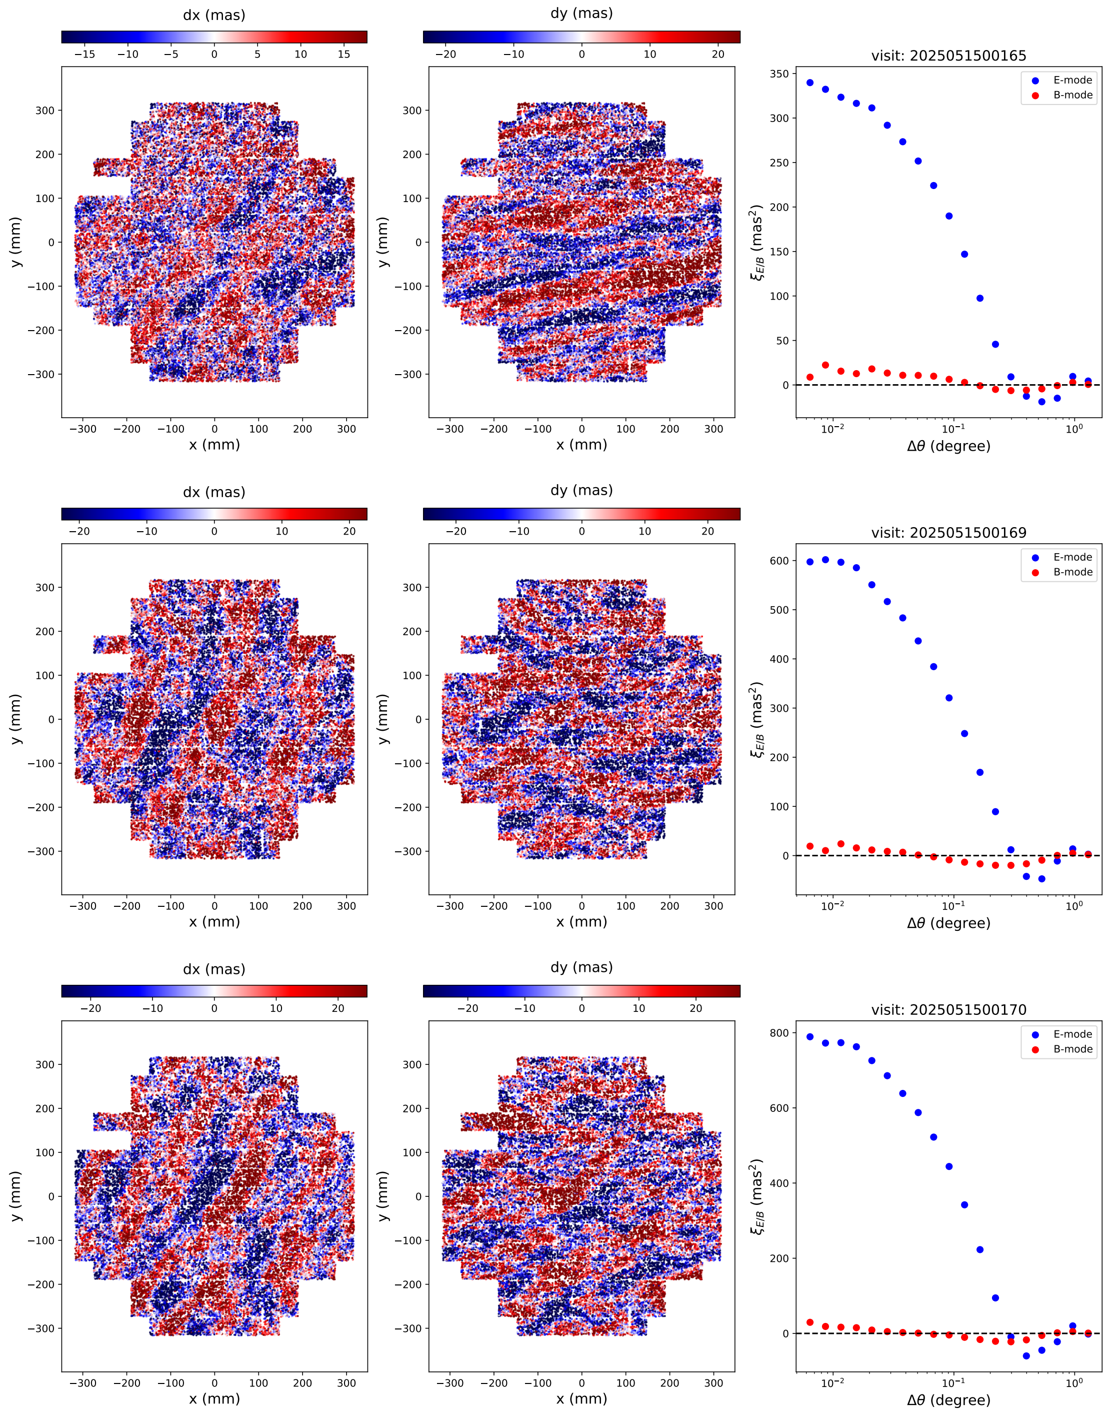

Figure 1. Astrometric residuals for three DP2 visits taken on the same night (2025051500165, 2025051500169, and 2025051500170). For each visit, the left and center panels show the x and y components of the residuals (in mas) as a function of position in the focal plane (in mm), and the right panel shows the corresponding E-mode (blue) and B-mode (red) correlation functions as a function of angular separation. The residual fields show the elongated stripe patterns expected from atmospheric turbulence, with an orientation that changes from one visit to the next, and their spatial correlations are dominated by the E-mode.

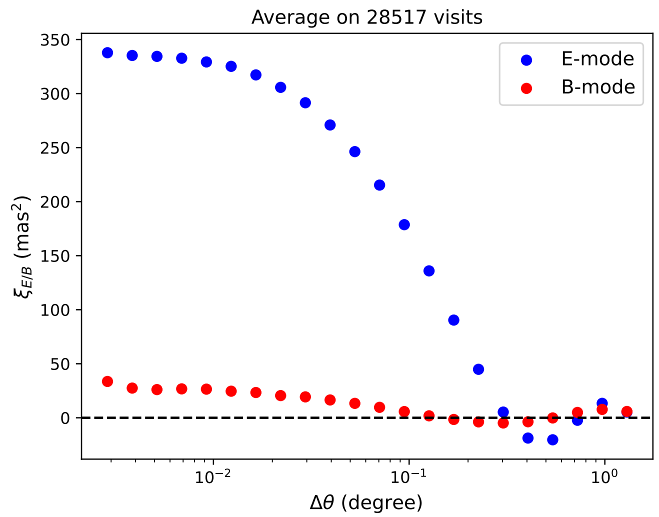

Figure 2. E-mode (blue) and B-mode (red) correlation functions of the astrometric residuals, averaged over the 28 517 DP2 visits used in this analysis, as a function of angular separation. The residual field is dominated by the E-mode, with a variance of $\sim$340 mas^2 at small separations, while the B-mode stays at the level of $\sim$30 mas^2.

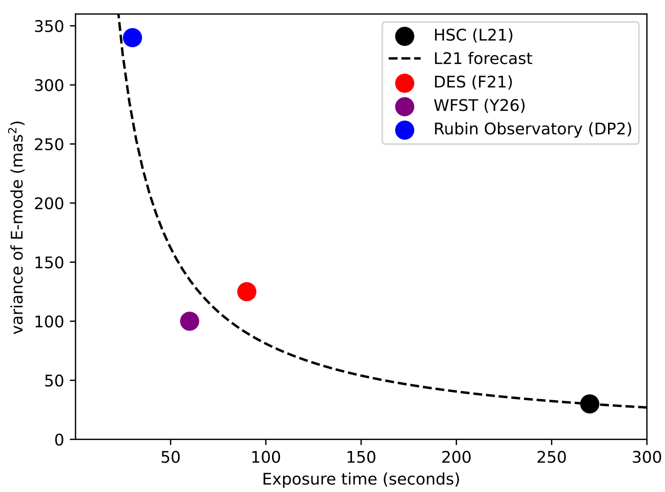

Figure 3. Variance of the E-mode of the astrometric residuals as a function of exposure time. The dashed line is the forecast made in L21, which rescales the HSC measurement (black) assuming that the variance scales as the inverse of the exposure time. The points show the values measured for DES (F21), WFST (Y26), and the Rubin Observatory with DP2 (this work). The DP2 value lies above the forecast, mostly because of the shorter exposure time of the Rubin Observatory.

where $\mathbf{K}$ is the covariance due to the turbulence field and $\mathbf{D} = \mathrm{diag}(\sigma_i^2)$ is the measurement noise. Because everything is Gaussian, the values $\mathbf{y}'$ of the field at any other set of positions $\mathbf{X}'$ are also Gaussian once the measurements are known, with mean and covariance (see also L21, Eqs. 8 and 10)

$$\mathbf{y}' = \boldsymbol{\Xi}(\mathbf{X}', \mathbf{X})\, \mathbf{C}^{-1}\, \mathbf{y}, \tag{3}$$

$$\mathrm{Cov}(\mathbf{y}') = \Xi(\mathbf{X}', \mathbf{X}') - \Xi(\mathbf{X}', \mathbf{X})\, \mathbf{C}^{-1}\, \Xi(\mathbf{X}, \mathbf{X}'), \tag{4}$$

where $\Xi(\mathbf{A}, \mathbf{B})_{ij} = \xi(\mathbf{A}_i - \mathbf{B}_j)$. Equation 3 is the GP prediction and Eq. 4 its uncertainty. The prediction is a linear combination of the measured values: each training source contributes with a weight that is large when it is strongly correlated with the target position and well measured, and the matrix $\mathbf{C}^{-1}$ accounts for the redundancy between neighboring training sources. Without measurement noise, the prediction at a training position is the measured value itself; with noise, the GP smooths the data by the amount that the correlation function and the noise allow. This estimator is the optimal linear interpolator of a Gaussian random field; it is known as kriging in geostatistics and as Wiener filtering in signal processing.

Two conditions have to be met in practice. First, $\mathbf{C}$ must be invertible for any set of positions, which requires $\mathbf{K}$ to be positive semi-definite, i.e., to have no negative eigenvalue. For a stationary kernel, this is guaranteed if and only if the Fourier transform of ξ is non-negative everywhere (Bochner's theorem; S. Bochner 1933). Analytic kernels such as the Gaussian or the von Kármán ones satisfy this condition by construction, but an arbitrary function, for instance the noisy measured two-point correlation function or a spline fit to it, does not in general; this is the reason why L21 chose a parametric kernel, and the central difficulty of using the measured correlation function directly, to which we come back in Sect. 3.3. Second, the computing cost. $\mathbf{C}$ is an $N \times N$ matrix, and the standard way of applying $\mathbf{C}^{-1}$ is a Cholesky factorization, which takes $O(N^3)$ operations and $O(N^2)$ memory; the prediction at M positions then requires the $M \times N$ cross-covariance matrix $\Xi(\mathbf{X}', \mathbf{X})$, i.e., $O(NM)$ operations and memory. With $N \sim 10^4$ sources per visit, this is what limits the size of the training sample in production, and what led G25 to split the DES field of view into patches.

Finally, the kernel ξ has to be chosen. The classical approach picks an analytic family, for instance a Gaussian or, for atmospheric turbulence, a von Kármán correlation function, with a few hyperparameters (an amplitude, a correlation length, and possibly an anisotropy), and finds their values by maximizing the likelihood of the training data. Each evaluation of the likelihood requires the inverse and the determinant of $\mathbf{C}$, hence a Cholesky factorization, and the optimizer needs many evaluations; for $N \sim 10^3$–10^4 this takes from minutes to hours per visit and per component (L21, their Fig. 6). The next two subsections describe two ways around this cost.

3.2. *L21 and DP2 implementation*

The von Kármán correlation function follows from the von Kármán spectrum of turbulence and has broader wings than a Gaussian, which L21 found to better describe the long-range correlations of the astrometric residuals. In its anisotropic form, the kernel reads

$$\xi(\mathbf{X}_i, \mathbf{X}_j) = \phi^2 \left[r_{ij}^2\right]^{5/6} K_{5/6}(2\pi\, r_{ij})\,, \qquad r_{ij}^2 = (\mathbf{X}_i - \mathbf{X}_j)^T \mathbf{L}^{-1} (\mathbf{X}_i - \mathbf{X}_j), \tag{5}$$

where ϕ^2 is the variance of the field, $K_{5/6}$ is the modified Bessel function of the second kind, and $\mathbf{L}$ is a symmetric positive 2×2 matrix that sets the correlation length in each direction,

$$\mathbf{L} = \begin{pmatrix} \cos\alpha & \sin\alpha \\ -\sin\alpha & \cos\alpha \end{pmatrix}^T \begin{pmatrix} \ell^2 & 0 \\ 0 & (q\ell)^2 \end{pmatrix} \begin{pmatrix} \cos\alpha & \sin\alpha \\ -\sin\alpha & \cos\alpha \end{pmatrix}. \tag{6}$$

The correlation length is ℓ along the direction α and $q\ell$ perpendicular to it, with $q \leq 1$; $q = 1$ is the isotropic case. L21 parametrize the pair (q, α) with two "shear" components (g_1, g_2), in exact analogy with the ellipticity of galaxies in weak lensing, so that $g_1 = g_2 = 0$ corresponds to an isotropic correlation and the kernel has four hyperparameters: ϕ, ℓ, g_1, and g_2.

Instead of maximizing the likelihood, L21 measure the two-dimensional two-point correlation function of the training sample on a grid of separations $(\Delta u, \Delta v)$ with `TreeCorr` (M. Jarvis et al. 2004; M. Jarvis 2015), whose tree algorithm scales as $O(N \log N)$, estimate the covariance of this map by bootstrap resampling of the sources, and fit Eq. 5 to the map by a weighted least-squares minimization,

$$\chi^2 = \left(\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}(\boldsymbol{\theta}) - k\right)^T \mathbf{W} \left(\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}(\boldsymbol{\theta}) - k\right), \tag{7}$$

where $\boldsymbol{\xi}$ is the measured map, $\hat{\boldsymbol{\xi}}(\boldsymbol{\theta})$ the analytic kernel evaluated on the same grid for the hyperparameters $\boldsymbol{\theta}$, $\mathbf{W}$ the inverse of the bootstrap covariance matrix, and k a constant. The constant is needed because the measured correlation function has negative lobes at large separations, since a correlation function measured after subtraction of the mean integrates to zero, and a positive kernel cannot reproduce them; k is not part of the kernel used for the interpolation. The $O(N^3)$ likelihood evaluations are thus replaced by the fit of a smooth function to a small map, and a single Cholesky factorization of $\mathbf{C}$ remains, for the prediction itself.

This method is implemented in the `treegp` package[2] and wrapped into a task of the LSST Science Pipelines (Rubin Observatory Science Pipelines Developers et al. 2026), which we now describe; Vera C. Rubin Observatory Team (2026) gives the production context. For each visit, the inputs are the astrometric residuals of Sect. 2 and their uncertainties, both provided per source by the `gbdes` fit, and the positions of the sources in the tangent plane of the visit. Among the sources of the visit, 10 000 are randomly selected, with a fixed random seed, to form the training sample; visits with fewer sources use them all. The remaining sources are not used by the fit and serve as a validation sample in Sect. 4. This cap on the size of the training sample follows from the $O(N^3)$ time and $O(N^2)$ memory of the dense factorization, which have to fit within the computing budget of a production job. The two components are then treated independently. For each of them, the two-point correlation function of the training sample is measured on a 21×21 grid of separations extending to $\pm 0.3°$, with inverse-variance weights, the anisotropic von Kármán kernel is fit to it following Eq. 7, and the matrix $\mathbf{C} = \mathbf{K} + \mathbf{D}$ of the training sample is built and factorized.

[2] https://github.com/PFLeget/treegp

The prediction is then done detector by detector, because the astrometric model is stored per detector. For each detector, the GP prediction (Eq. 3) is evaluated on a regular 30×30 grid of tangent-plane positions covering the detector plus a margin of $0.1°$; the predicted displacement is fit by a two-dimensional spline that is appended to the astrometric model of the detector (as an AST `SplineMap`; Vera C. Rubin Observatory Team 2026), so that the position of any source on the detector is corrected through the WCS. The GP uncertainty (Eq. 4) is not computed in production.

The DP2 production revealed three limitations of this approach. First, the shape of the kernel is a choice. A single sheared von Kármán profile describes the dominant scale and orientation of the correlations, but the measured two-point correlation function frequently shows more structure, such as several scales or orientations and negative lobes; this is visible in the residual maps of L21 (their Fig. 8). Whatever the kernel does not describe is left in the corrected residuals. Second, the hyperparameter fit is a non-linear minimization that starts from an initial guess; on some visits it fails to converge, or converges to a poor solution, for instance when the correlation function is noisy because of a low number of sources. The 28 517 visits of Sect. 2 are those for which this fit succeeded, and the failure rate is discussed in Sect. 4.1. Third, the run time is dominated by the evaluation of the kernel itself: Eq. 5 requires a Bessel function for each pair of positions, i.e., 10^8 evaluations for the training matrix, and as many again as there are grid points times training sources for each detector at prediction time.

G25 address the first two limitations by construction: the measured two-point correlation function itself is the kernel, so there is no functional form and nothing to fit, and no Bessel function either. The price is the loss of the positive semi-definiteness guarantee of Sect. 3.1, which G25 recover with an $O(N^3)$ singular value decomposition. TURBO-GP starts from G25 and removes this last obstacle.

3.3. *TUrbulence Removal By fourier-Optimized Gaussian Process*

The kernel of TURBO-GP is the one of G25, with one addition, and is built in three steps for each component of the residual field. First, the two-dimensional two-point correlation function of the training sample is measured with `TreeCorr` on a square grid of separations, with inverse-variance weights; the configuration used in this paper has $20''$ pixels and extends to $\pm 0.3°$. The map is rebinned by 2×2 so that the zero separation falls at the center of a pixel. Second, the map is cleaned of its measurement noise. It is multiplied by an apodization window (F. J. Harris 1978) that goes smoothly to zero at a chosen radius, here a Hann window reaching zero at $0.278°$ (G25 use a Blackman–Harris window), which prevents the sharp edge of the measured square from ringing in Fourier space. Its Fourier transform is taken, and every mode whose power is below 2.5 times the noise level, estimated from the scatter of the high-frequency half of the spectrum, is set to zero; the map is then transformed back. The threshold removes the noise of the measurement and leaves a spectrum that is non-negative and made of a small number of Fourier modes, two properties that are central to what follows. Third, and this is the addition with respect to G25, the apodization window is made elliptical and matched to the anisotropy of the correlation function, so that an elongated correlation pattern (Fig. 1) is not truncated along its long axis: the map is first cleaned with an isotropic window, its anisotropy is measured with adaptive weighted second moments (C. Hirata & U. Seljak 2003) and expressed with the shear parametrization (g_1, g_2) of L21, and the raw map is cleaned again with the corresponding elliptical window. The cleaned map, linearly interpolated between its pixels, is the kernel. There are no hyperparameters and nothing to fit. Unlike G25, we do not subtract a third-order polynomial before the GP, since the per-visit part of the astrometric model of Sect. 2 already absorbs the large-scale distortions**[TODO: PFL: confirm this justification]**, we do not include the cross-correlation between the two components, and, as explained below, we do not divide the field of view into patches.

The difficulty with this kernel is the one anticipated in Sect. 3.1. The cleaned map has a non-negative spectrum, so it is a valid kernel as a periodic function, with a period equal to the size of the measured square (0.6° here); but the field of view is much larger than this period, and the kernel actually used sets the correlation to zero beyond the measured square. This truncation is what breaks the positive semi-definiteness: in Fourier space it convolves the spectrum with an oscillating function and creates negative lobes, so that the matrix $\mathbf{K}$ between arbitrary source positions acquires negative eigenvalues whose magnitude is a non-negligible fraction of the variance of the field. The Cholesky factorization of $\mathbf{C}$ then fails whenever some sources have small measurement uncertainties, which is exactly the case of the bright sources that constrain the field the most. G25 repair $\mathbf{K}$ with a singular value decomposition whose small and negative singular values are clipped to a positive floor, at an $O(N^3)$ cost for each visit and each component.

TURBO-GP takes a different route, which makes the kernel positive semi-definite by construction and never forms the matrix $\mathbf{K}$. The observation is that the cleaned map is exactly a finite sum of Fourier modes on its own grid: it is band-limited, so its values anywhere are known exactly from its values on a grid. We therefore represent the covariance between sources through a fine regular grid of the tangent plane with a spacing h equal to half the pixel of the correlation function ($10''$ here), covering the field of view enlarged by the extent of the kernel. Each source is spread on the four grid nodes that surround it with bilinear weights, which defines a sparse $N \times N_g$ matrix $\mathbf{W}$ with four non-negative entries per row summing to one, where N_g is the number of grid nodes. The covariance between two grid nodes depends only on their separation, so the covariance between all nodes is a convolution by the kernel image, which we write as a matrix $\mathbf{G}$ and compute with fast Fourier transforms on a grid zero-padded to more than the size of the field plus the extent of the kernel, so that the periodic wrap of the discrete Fourier transform never connects two sources. The covariance between sources is then

$$\mathbf{K} = \mathbf{W}\,\mathbf{G}\,\mathbf{W}^T, \qquad \mathbf{G} = \frac{1}{N_g}\,\mathbf{F}^\dagger\,\mathrm{diag}(\hat{P}_k)\,\mathbf{F}, \tag{8}$$

where $\mathbf{F}$ is the two-dimensional discrete Fourier transform on the padded grid and $\hat{P}_k$ the Fourier coefficients of the kernel image, which are real because the image is symmetric. Zero-padding the kernel image is a truncation again, and a few of the $\hat{P}_k$ are slightly negative; we clip them to zero, once, at the cost of one Fourier transform. Then, for any vector $\mathbf{v}$,

$$\mathbf{v}^T\mathbf{K}\,\mathbf{v} = \frac{1}{N_g}\sum_k \hat{P}_k \left|(\mathbf{F}\,\mathbf{W}^T\mathbf{v})_k\right|^2 \geq 0, \tag{9}$$

which is the statement that $\mathbf{K}$ is positive semi-definite, for any configuration of sources and not only for the training sample of a given visit. The matrix $\mathbf{C} = \mathbf{K} + \mathbf{D}$ is therefore positive definite as soon as every source has a non-zero uncertainty, and a failure of the linear algebra is not repaired but impossible. Equation 9 is nothing else than Bochner's theorem applied to the padded grid: the effective kernel is the cleaned correlation function smoothed by the bilinear interpolation at the scale h, a second-order effect, with the non-negative spectrum $\hat{P}_k$. The per-visit $O(N^3)$ decomposition of G25 is not accelerated but eliminated.

The linear system $\mathbf{C}\,\boldsymbol{\alpha} = \mathbf{y}$ of Eq. 3 is solved without ever storing $\mathbf{C}$, with the conjugate gradient method (M. R. Hestenes & E. Stiefel 1952) and a diagonal (Jacobi) preconditioner. Each iteration applies $\mathbf{C}$ to a vector, which by Eq. 8 is one spreading on the grid, one pair of fast Fourier transforms, one interpolation back to the sources, and one multiplication by $\mathbf{D}$: $O(N + N_g \log N_g)$ operations and $O(N + N_g)$ memory, instead of $O(N^3)$ and $O(N^2)$. The iterations stop at a fixed relative tolerance, so that the solution is deterministic; their number grows with the ratio of the variance of the field to the measurement noise, and is of the order of a hundred on the DP2 data.

The prediction is also reorganized. Once $\boldsymbol{\alpha} = \mathbf{C}^{-1}\mathbf{y}$ is known, the GP prediction of Eq. 3 at any position $\mathbf{X}$ is a field,

$$\hat{y}(\mathbf{X}) = \sum_{j=1}^{N} \alpha_j\,\xi(\mathbf{X} - \mathbf{X}_j), \tag{10}$$

that is, the convolution of the kernel with the weights α_j placed at the training positions. With the grid representation, this field is computed once per component, as a single convolution of the spread weights $\mathbf{W}^T\boldsymbol{\alpha}$ by the kernel image, and stored as an image of the focal plane. Each of the per-detector predictions of Sect. 3.2 is then a bilinear read-off of this image, with a cost independent of the number of training sources, in place of the $M \times N$ cross-covariance matrix of the dense approach, which dominated the run time of G25. Positions farther than the extent of the kernel from every training source correctly return zero. As in DP2, the GP uncertainty is not computed in production.

Since both the solve and the prediction now scale essentially linearly with the number of sources, there is no need to divide the field of view into patches, which G25 introduced only to keep the dense linear algebra affordable: TURBO-GP fits one GP per component over the whole focal plane, with no buffer zones and no discontinuities at patch boundaries, and the number of training sources is no longer limited by the linear algebra. For the comparison presented in this paper, however, TURBO-GP is run exactly like the L21 model of DP2: the same 10 000 randomly selected training sources, with the same random seed, the same per-source uncertainties in $\mathbf{D}$, and the same per-detector prediction on a 30×30 grid followed by the same spline insertion in the WCS. Only the kernel and the way the linear algebra is solved differ, so that the differences observed in Sect. 4 can be attributed to the model itself. The gain from larger training samples, which TURBO-GP makes affordable, is left for future work**[TODO: PFL: or point to the results section if it is tested there]**.

None of the numerical ingredients of TURBO-GP is new in itself. Representing a GP covariance through a regular grid with sparse local interpolation and a fast Fourier transform convolution is known in the machine-learning literature as structured kernel interpolation (KISS-GP; A. G. Wilson & H. Nickisch 2015), and the spread–transform–interpolate sequence is the structure of non-uniform fast Fourier transforms. What is specific to TURBO-GP is the observation that the empirical kernel of G25 is band-limited on its own measurement grid, so that the grid representation, usually an approximation, is here essentially exact, and that positive semi-definiteness can be imposed once in Fourier space instead of being repaired for each visit. This is what allows the same kernel to be used over a full LSSTCam focal plane within the computing budget of the Rubin production, as quantified in Sect. 4.

4. RESULTS

4.1. *Time and memory performance on DP2*

We re-ran L21 and TURBO-GP side by side on DP2 in order to compare them. We used the Batch Production System (T. Jenness et al. 2022) at the SLAC Shared Scientific Data Facility (S3DF), the facility used to produce DP2 (Vera C. Rubin Observatory Team 2026). Both L21 and TURBO-GP are implemented in the `treegp` package and wrapped in the LSST Science Pipelines (Rubin Observatory Science Pipelines Developers et al. 2026), in the `drp_tasks` package, through the `gaussianProcessesTurbulenceFit` task, so that comparing the two methods only requires changing one configuration parameter when the task is executed. For the reasons given in Sect. 2, a visit may be processed several times when it overlaps several HEALPix regions; the numbers quoted below are per job, i.e., per (visit, HEALPix region) pair.

The first improvement is on the error rate, defined as the number of jobs that did not complete over the total number of jobs. On DP2, L21 shows an error rate of 5%, while TURBO-GP shows an error rate of 1%. This is explained mostly by failures of the hyperparameter minimization in L21 (Sect. 3.2), which lead to a kernel for which the Cholesky decomposition fails. Since this step is entirely removed in TURBO-GP, these failures no longer occur.

We can then compare the computing time of the two methods; the distributions of the wall time per job are shown in Figure 4. There is an improvement of almost two orders of magnitude from L21 to TURBO-GP, which brings the total wall time for DP2 from ∼2.7 years down to ∼2 weeks. The difference is explained by several factors. At training, there is no fit anymore: instead of computing the two-point correlation function several hundred times, as required by the bootstrap estimate of its covariance matrix, and then fitting the hyperparameters, the two-point correlation function is computed only once. At inference, there are two factors. First, because TURBO-GP is empirical, it only needs the tabulated two-point correlation function and no longer calls the von Kármán function, which is slow because of its Bessel function. Second, there is no Cholesky decomposition anymore: the $O(N^3)$ factorization is replaced by the conjugate-gradient solve of Eq. 8, whose iterations cost $O(N + N_g \log N_g)$, and the $O(MN)$ multiplication by the cross-covariance matrix is replaced by the single convolution of Eq. 10, followed by a read-off whose cost is $O(M)$.

Figure 5 presents the maximum memory usage during job execution. The memory usage is also significantly improved by TURBO-GP, which follows directly from never building the $N \times N$ covariance matrix and its factorization, nor the $M \times N$ cross-covariance matrix. However, there is a long tail at high memory usage where TURBO-GP and L21 match. This is due to a design choice in the task. At the time of writing, `gaussianProcessesTurbulenceFit` loads in memory a table that contains the sources of all the visits of a given HEALPix region, and then selects the sources of the visit being processed; HEALPix regions with more visits therefore require more memory just to select one visit. When the task was developed for DP2, this choice minimized data manipulation, but it will not scale to future data releases, and it will be changed.

Improving the memory usage and the computing time normally comes at the cost of a less accurate method. However, we show in the following subsections that, because TURBO-GP uses the empirical two-point correlation function at its root, it also improves how well atmospheric turbulence is described for astrometry.

4.2. *On single visits of Rubin Observatory*

We show the results on a single visit (visit 2025121400911) in Figure 6 for the L21 model used in the DP2 production and in Figure 7 for TURBO-GP. While both L21 and TURBO-GP describe the overall features of the residual field well, TURBO-GP visually describes the small-scale features better. This can be seen in the correlation functions of the corrected residuals: the overall astrometric residuals after GP modeling are smaller for TURBO-GP than for L21, by about 10% in the variance of the

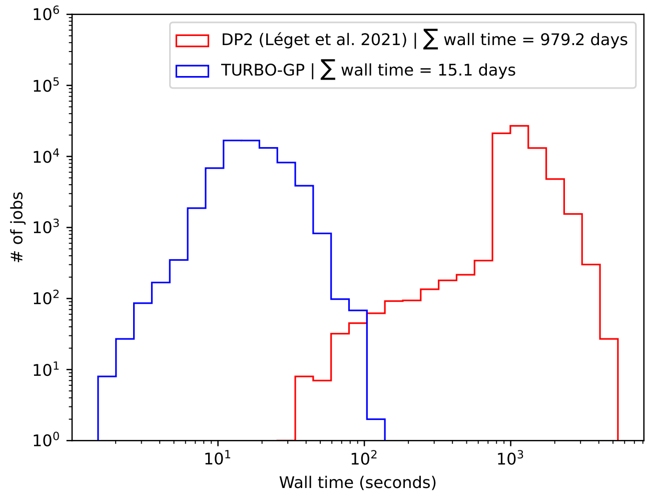

Figure 4. Distribution of the wall time per job of the `gaussianProcessesTurbulenceFit` task on DP2, with one job per (visit, HEALPix region) pair, for the L21 model used in the DP2 production (red) and for TURBO-GP (blue). The legend gives the wall time summed over all jobs: 979 days for L21 and 15 days for TURBO-GP.

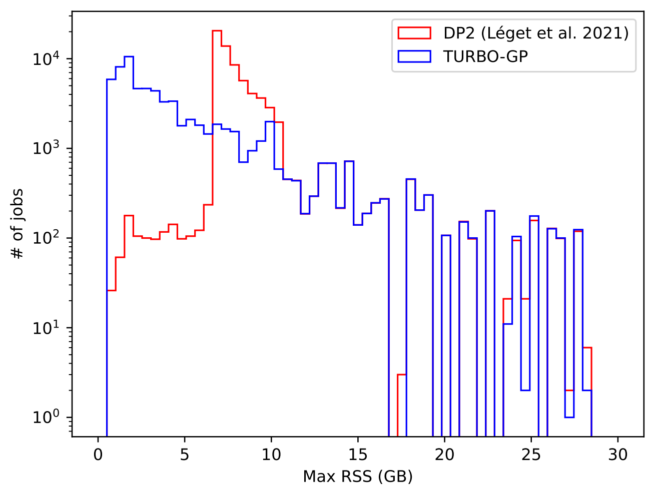

Figure 5. Distribution of the maximum resident memory (max RSS) per job of the `gaussianProcessesTurbulenceFit` task on DP2, for the L21 model used in the DP2 production (red) and for TURBO-GP (blue). The bulk of the L21 jobs peaks at ∼7 GB, while the TURBO-GP jobs peak at ∼1–2 GB. The common tail up to ∼28 GB is due to the loading of the source table of the whole HEALPix region by the task (see text).

E-mode. This is an example on a single visit, but we show in the next subsection that it generalizes to the full DP2 sample.

4.3. *Global results on DP2*

For the 28 517 visits in common between the DP2 L21 run and TURBO-GP, we compute the E- and B-mode correlation functions of the astrometric residuals cor-

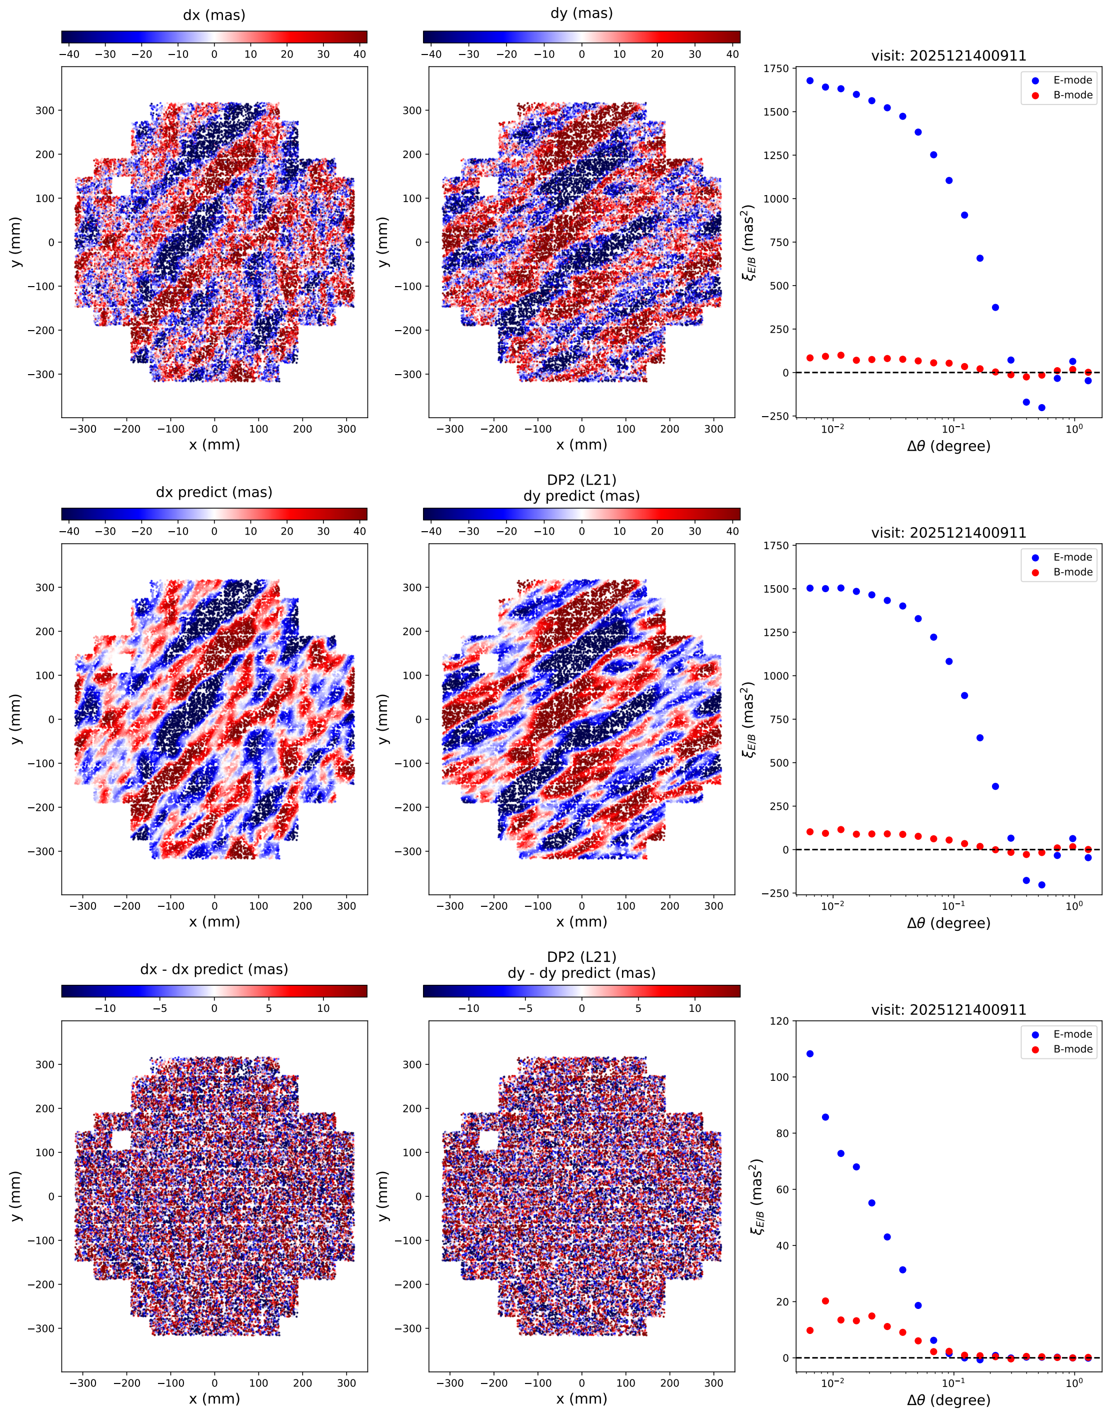

Figure 6. Atmospheric turbulence modeling of visit 2025121400911 with the L21 model used in the DP2 production. Top row: x (left) and y (center) components of the astrometric residuals (in mas) as a function of position in the focal plane (in mm), and their E-mode (blue) and B-mode (red) correlation functions as a function of angular separation (right); this visit is strongly affected by turbulence, with an E-mode variance of $\sim$1700 mas^2, about five times the DP2 average (Fig. 2). Middle row: GP prediction at the positions of the sources, with the same color scale, and the correlation functions of the prediction. Bottom row: residuals after subtraction of the GP prediction, with a color scale reduced by a factor of $\sim$3, and their correlation functions; note the change of vertical scale, the E-mode variance at the smallest separations going from $\sim$1700 to $\sim$110 mas^2.

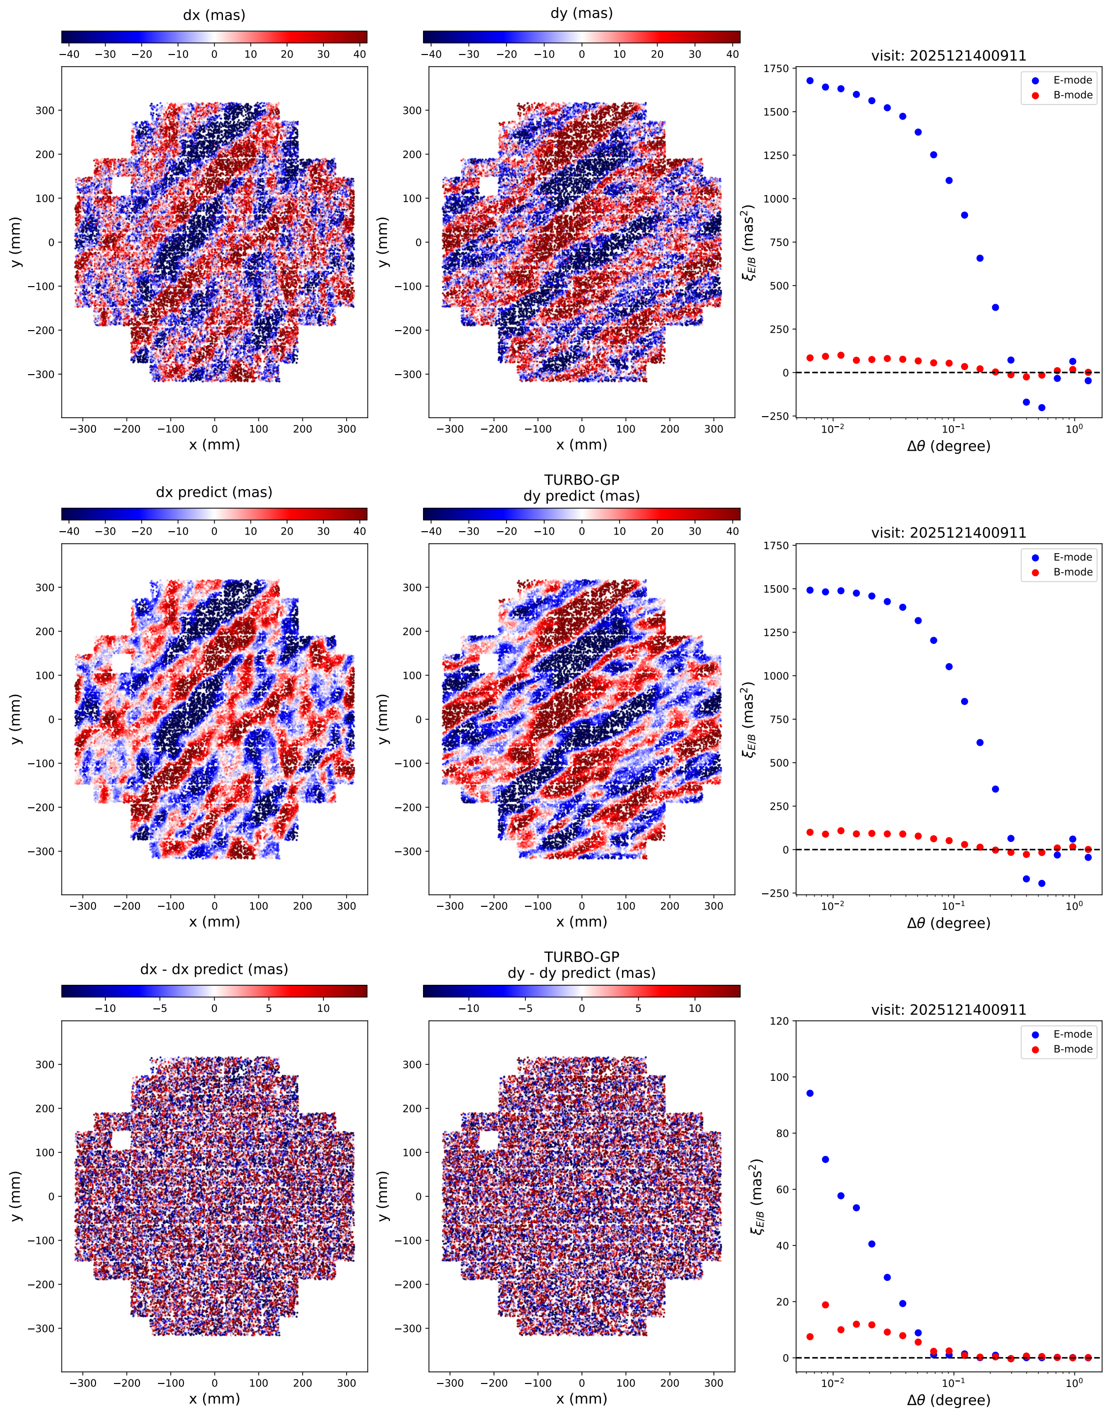

Figure 7. Same as Fig. 6 for TURBO-GP, on the same visit and with the same training sample. The prediction (middle row) follows the small-scale features of the residual field more closely, and the E-mode variance of the corrected residuals at the smallest separations (bottom right) is $\sim$95 mas^2, about 10% lower than with L21.

rected with the respective GP model, and we average these correlation functions over the 28 517 visits. The results are presented in Figure 8. Both L21 and TURBO-GP remove $\sim$90% of the variance introduced by the atmosphere (Fig. 2), but the residual

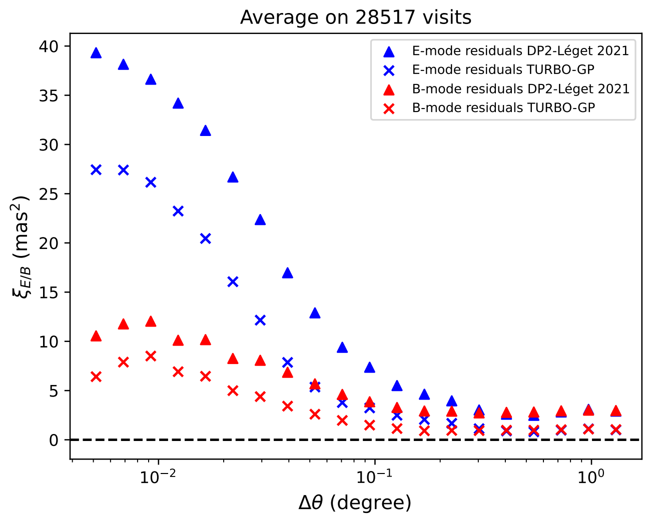

Figure 8. E-mode (blue) and B-mode (red) correlation functions of the astrometric residuals after correction of the atmospheric turbulence, averaged over the 28 517 DP2 visits, for the L21 model used in the DP2 production (triangles) and for TURBO-GP (crosses). Compared with the $\sim$340 mas^2 of E-mode variance before correction (Fig. 2), the residual E-mode variance at the smallest separations is $\sim$39 mas^2 for L21 and $\sim$27 mas^2 for TURBO-GP; the B-mode goes from $\sim$30 mas^2 to $\sim$11 and $\sim$7 mas^2, respectively. Beyond $\sim 0.3°$, both methods leave a floor of a few mas^2 for L21 and $\sim$1 mas^2 for TURBO-GP.

variance of TURBO-GP is 30% lower in E-mode and 40% lower in B-mode than that of L21. So even if L21 was already doing a good job at keeping atmospheric turbulence under control, TURBO-GP improves the description of the atmosphere, at a lower cost in time and memory, and with a description of the spatial correlation that is measured directly from the data rather than imposed by a parametric kernel.

There is still some remaining spatial correlation at angular scales well below the size of a CCD itself, and we discuss one potential candidate in Sect. 4.6. But before that, we discuss the correlation of the astrometric residuals with image quality.

4.4. *Discussion*

We use this discussion to examine the correlation of the astrometric residuals with image quality, and what remains to be modeled in the spatially structured astrometric distortions.

4.5. *Correlation with image quality*

Throughout the period covered by the DP2 data set, a major focus of the Rubin Observatory team was to bring the image quality (IQ) under control, in particular through the Active Optics System (AOS). The complication with IQ is that it is unclear whether the PSF is driven by optical misalignments or by atmospheric turbulence, and breaking the degeneracy between the two is a well-known challenge.

The advantage of the analysis presented here with the astrometric residuals is that they are unambiguously driven by the atmosphere alone: the optical distortions are taken into account by the astrometric model, and the residuals are mostly E-mode,

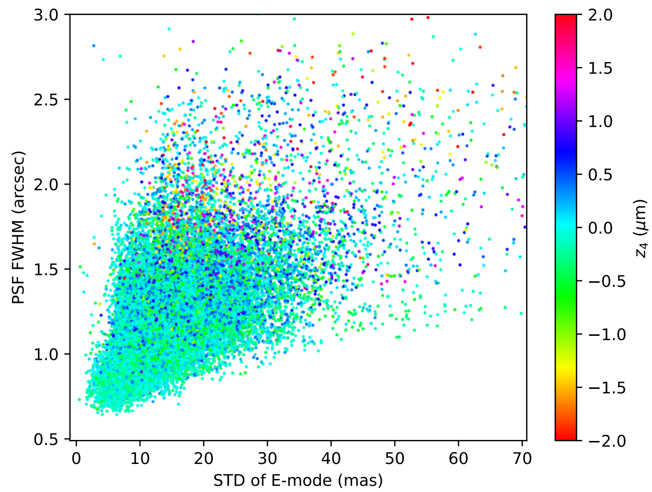

Figure 9. Median PSF FWHM (in arcsec) of each DP2 visit as a function of the standard deviation of the E-mode of its astrometric residuals (in mas), colored by the median defocus Zernike coefficient z_4 (in μm) measured by the AOS. The PSF FWHM correlates with the E-mode of the astrometric residuals, i.e., with the strength of the atmospheric turbulence, independently of the focus.

which indicates that the astrometric shifts follow the gradient of the spatial variations of the refractive index integrated along the line of sight (Sect. 1). In short, they are a clean signature of what is going on in the atmosphere.

As far as we know, a cross-correlation between astrometric residuals and image quality has never been observed for a ground-based optical wide-field telescope. To do so, we compute for every visit the standard deviation of the E-mode of the astrometric residuals, the median FWHM of the PSF, and the median value of z_4 over the four wavefront measurements of the AOS. In the Noll ordering of the Zernike polynomials, z_4 is the defocus term; we look only at this coefficient because it is the one that corresponds to the focus of the telescope and is expected to correlate directly with the PSF FWHM. Figure 9 shows the PSF FWHM as a function of the standard deviation of the E-mode (colored by the median z_4), and Figure 10 shows the PSF FWHM as a function of z_4 (colored by the standard deviation of the E-mode).

Figure 9 shows a clear correlation between the PSF FWHM and the standard deviation of the E-mode of the astrometric residuals. This shows that, as expected, a major driver of the PSF FWHM is atmospheric turbulence. Figure 10 shows the expected parabola between the PSF FWHM and the focus, but a large part of the scatter around it correlates with the standard deviation of the E-mode, which is evidence that this scatter is driven by atmospheric turbulence.

4.6. *On static defects from on-sky measurements*

A. A. Plazas et al. (2014a) and L21 showed that static distortions introduced by CCD features such as tree rings can be observed directly from on-sky measurements,

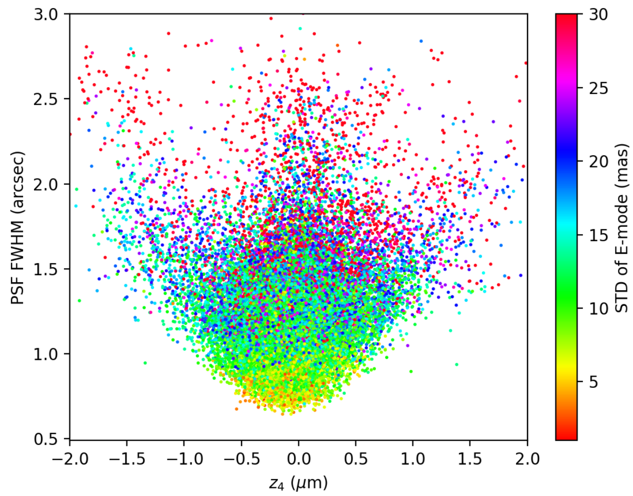

Figure 10. Median PSF FWHM (in arcsec) of each DP2 visit as a function of the median defocus Zernike coefficient z_4 (in μm) measured by the AOS, colored by the standard deviation of the E-mode of the astrometric residuals (in mas). The parabolic dependence of the PSF FWHM on the focus is visible, and the scatter above the parabola is populated by the visits with a high E-mode, i.e., a strong atmospheric turbulence.

and in the same spirit we can do the same here. After correcting for atmospheric turbulence with TURBO-GP, we bin every sensor of LSSTCam in super-pixels of 10×10 pixels and compute a clipped median of the corrected residuals across all the visits of DP2. Results for one CCD of each kind (ITL and E2V) are displayed in Figure 11 (detector 1, ITL) and Figure 12 (detector 97, E2V).

Both detectors exhibit the expected features. Both the ITL and the E2V sensors show tree rings, and in addition, as was measured with artificial stars by J. H. Esteves et al. (2023), the imprint of the mechanical support of the ITL sensor is clearly visible. The ITL sensor exhibits typical systematic shifts of 4 mas, while the E2V sensor exhibits shifts of 1–2 mas, which can explain part of the spatial correlation remaining after the atmospheric turbulence is properly modeled with TURBO-GP (Fig. 8).

At the time of writing, these effects are not yet corrected, as was done for example for DES (A. A. Plazas et al. 2014a,b). The way to do it would in principle be to derive the correction from the flat fields, as was done for DES. But, as we can see, it can also be derived directly from on-sky measurements once atmospheric turbulence has been corrected. We leave this for further analysis, but this is the clear remaining path to improve the astrometric precision of LSSTCam for the Rubin Observatory.

5. CONCLUSION

We have used the 28 517 visits of the Rubin Observatory Data Preview 2 with a successful atmospheric turbulence fit to characterize how atmospheric turbulence affects the astrometry of LSSTCam. Once the optical distortions are fit by the astrometric model, the residual displacement field is dominated by the E-mode, with a variance of

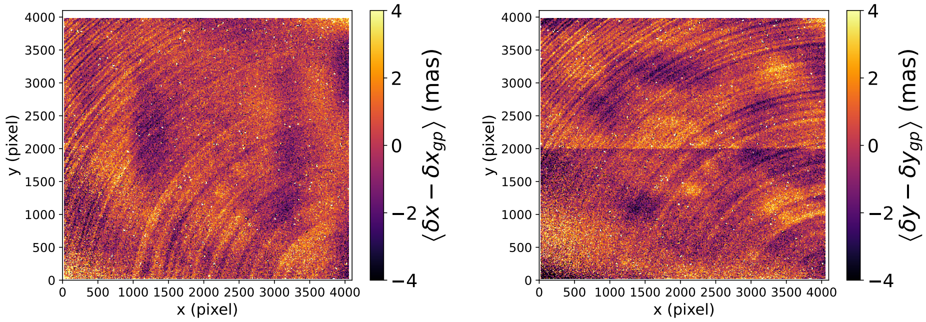

Figure 11. Static astrometric residuals of detector 1 (ITL) in its pixel coordinates: clipped median, across all the DP2 visits, of the astrometric residuals corrected for atmospheric turbulence with TURBO-GP, in super-pixels of 10×10 pixels, for the x (left) and y (right) components (in mas). Concentric tree rings centered on the upper-left corner of the sensor are visible, together with oblique stripes that we associate with the mechanical support of the sensor (see text) and a discontinuity at the mid-line of the sensor ($y \simeq 2000$ pixels), which separates its two rows of amplifiers. The typical amplitude of these features is 4 mas.

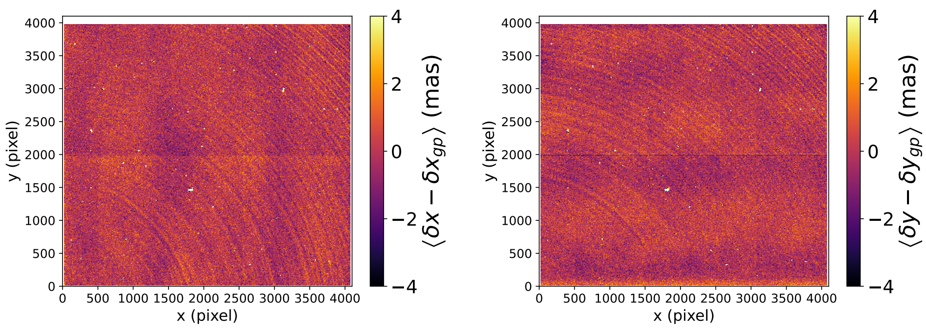

Figure 12. Same as Fig. 11 for detector 97 (E2V). Tree rings centered on the upper-right corner of the sensor are visible, with an amplitude of 1–2 mas, i.e., a factor of 2–4 smaller than on the ITL sensor, together with the discontinuity at the mid-line of the sensor.

$\sim$340 mas^2, ten times higher than what was observed with HSC and three times higher than with DES. This is slightly above the 270 mas^2 that L21 forecast by rescaling the HSC measurement to the 30 s exposure time of the Rubin Observatory, and it makes the correction of atmospheric turbulence more important for the Rubin Observatory than for any previous wide-field survey.

The DP2 production corrected this field with the Gaussian process of L21, with a sheared von Kármán kernel whose hyperparameters are fit to the measured two-point correlation function. This method removes $\sim$90% of the variance introduced by the atmosphere, but the production revealed its limitations: the shape of the kernel is a choice that does not capture all the structure of the measured correlation function, the hyperparameter fit fails on a fraction of the visits, and the evaluation of the von Kármán kernel is slow. Following G25, TURBO-GP uses the measured two-point correlation function directly as the kernel, so that there are no hyperparameters and nothing to fit. The new element is the way the linear algebra is solved: the covariance is represented on a regular grid of the focal plane, the kernel is applied by fast Fourier transform, positive semi-definiteness is imposed once in Fourier space rather than repaired for each visit with an $O(N^3)$ decomposition, the linear system is solved by conjugate gradient, and the prediction is computed once as a field over the whole focal plane. As a result, there is no need to divide the field of view into patches, and one Gaussian process per component is fit over the whole LSSTCam focal plane. Both methods are implemented in `treegp` and in the LSST Science Pipelines, and switching from one to the other only requires changing one configuration parameter.

On DP2, TURBO-GP reduces the error rate of the turbulence-fitting task from 5% to 1%, the total wall time from ∼2.7 years to ∼2 weeks, and the typical memory usage per job from ∼7 GB to ∼1–2 GB. This improvement in cost does not come with a loss of accuracy: on the same training samples, the astrometric residuals corrected with TURBO-GP have, averaged over DP2, an E-mode variance 30% lower and a B-mode variance 40% lower than those corrected with the L21 model of the DP2 production. TURBO-GP is therefore the method we intend to use for future data releases, and the larger training samples it makes affordable remain to be explored.

The astrometric residuals corrected in this way also turn out to be a clean probe of the atmosphere. Across DP2, the standard deviation of the E-mode of the astrometric residuals correlates with the PSF FWHM, and it explains a large part of the scatter of the PSF FWHM around its expected parabolic dependence on the focus; to our knowledge, this is the first time such a correlation is observed for a ground-based optical wide-field telescope. Finally, once atmospheric turbulence is removed, the residuals averaged over all the visits of DP2 in the pixel coordinates of each sensor reveal the static distortions of LSSTCam: tree rings on both the ITL and the E2V sensors, and the imprint of the mechanical support of the ITL sensors, with typical amplitudes of 4 mas for ITL and 1–2 mas for E2V. These static distortions are not yet corrected in the Rubin astrometric model and account for part of the spatial correlation left at small scales; deriving their correction, from flat fields or directly from these on-sky measurements, is the clear next step to improve the astrometric precision of the Rubin Observatory.

ACKNOWLEDGMENTS

This material is based upon work supported in part by the National Science Foundation through Cooperative Agreements AST-1258333 and AST-2241526 and Cooperative Support Agreements AST-1202910 and AST-2211468 managed by the Association of Universities for Research in Astronomy (AURA), and the Department of Energy under Contract No. DE-AC02-76SF00515 with the SLAC National Accelerator Laboratory managed by Stanford University. Additional Rubin Observatory funding comes from private donations, grants to universities, and in-kind support from LSST-DA Institutional Members.

Facility: Rubin

Software: LSST Science Pipelines

REFERENCES

Bernardinelli, P. H., Bernstein, G. M., Sako, M., et al. 2020, ApJS, 247, 32, doi: 10.3847/1538-4365/ab6bd8

Bernstein, G. M. 2022, gbdes: DECam instrumental signature fitting and processing programs„ Astrophysics Source Code Library, record ascl:2210.011 http://ascl.net/2210.011

Bernstein, G. M., Armstrong, R., Plazas, A. A., et al. 2017, PASP, 129, 074503, doi: 10.1088/1538-3873/aa6c55

Bochner, S. 1933, Mathematische Annalen, 108, 378, doi: 10.1007/BF01452844

Esteves, J. H., Utsumi, Y., Snyder, A., et al. 2023, PASP, 135, 115003, doi: 10.1088/1538-3873/ad0a73

Ferguson, P. S., Rykoff, E. S., Carlin, J. L., Saunders, C., & Parejko, J. K. 2025, The Monster: A reference catalog with synthetic ugrizy-band fluxes for the Vera C. Rubin observatory, Data Management Technical Note DMTN-277, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2583688

Fortino, W. F., Bernstein, G. M., Bernardinelli, P. H., et al. 2021, AJ, 162, 106, doi: 10.3847/1538-3881/ac0722

Gomes, D. C. H., Bernstein, G. M., & Hébert, C.-A. 2025, AJ, 170, 361, doi: 10.3847/1538-3881/ae1a7b

Harris, F. J. 1978, Proceedings of the IEEE, 66, 51, doi: 10.1109/PROC.1978.10837

Hestenes, M. R., & Stiefel, E. 1952, Journal of Research of the National Bureau of Standards, 49, 409, doi: 10.6028/jres.049.044

Hirata, C., & Seljak, U. 2003, MNRAS, 343, 459, doi: 10.1046/j.1365-8711.2003.06683.x

Ivezić, Ž., & The LSST Science Collaboration. 2018, LSST Science Requirements Document, Project Controlled Document LPM-17, NSF-DOE Vera C. Rubin Observatory. https://ls.st/LPM-17

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

Jarvis, M. 2015, TreeCorr: Two-point correlation functions„ Astrophysics Source Code Library, record ascl:1508.007 http://ascl.net/1508.007

Jarvis, M., Bernstein, G., & Jain, B. 2004, MNRAS, 352, 338, doi: 10.1111/j.1365-2966.2004.07926.x

Jenness, T., Bosch, J. F., Salnikov, A., et al. 2022, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 12189, Software and Cyberinfrastructure for Astronomy VII, 1218911, doi: 10.1117/12.2629569

Léget, P. F., Astier, P., Regnault, N., et al. 2021, A&A, 650, A81, doi: 10.1051/0004-6361/202140463

NSF-DOE Vera C. Rubin Observatory. 2026, Legacy Survey of Space and Time Data Preview 2 [Data set], NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/3382528

Park, H. Y., Nomerotski, A., & Tsybychev, D. 2017, Journal of Instrumentation, 12, C05015, doi: 10.1088/1748-0221/12/05/C05015

Plazas, A. A., & Bernstein, G. 2012, PASP, 124, 1113, doi: 10.1086/668294

Plazas, A. A., Bernstein, G. M., & Sheldon, E. S. 2014a, Journal of Instrumentation, 9, C04001, doi: 10.1088/1748-0221/9/04/C04001

Plazas, A. A., Bernstein, G. M., & Sheldon, E. S. 2014b, PASP, 126, 750, doi: 10.1086/677682

Rasmussen, C. E., & Williams, C. K. I. 2006, Gaussian Processes for Machine Learning (Cambridge, MA: MIT Press)

Rubin Observatory Science Pipelines Developers, AlSayyad, Y., Axelrod, T. S., et al. 2026, The LSST Science Pipelines Software: Optical Survey Pipeline Reduction and Analysis Environment, Project Science Technical Note PSTN-019, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2570545

Saunders, C. 2024, Astrometric Calibration in the LSST Pipeline, Data Management Technical Note DMTN-266, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/2583846

SLAC National Accelerator Laboratory, & NSF-DOE Vera C. Rubin Observatory. 2025, The LSST Camera (LSSTCam), SLAC National Accelerator Laboratory (SLAC), Menlo Park, CA (United States), doi: 10.71929/rubin/2571927

Vera C. Rubin Observatory Team. 2026, The Vera C. Rubin Observatory Data Preview 2, Technical Note RTN-115, NSF-DOE Vera C. Rubin Observatory, doi: 10.71929/rubin/3377440

Wilson, A. G., & Nickisch, H. 2015, in Proceedings of Machine Learning Research, Vol. 37, Proceedings of the 32nd International Conference on Machine Learning (PMLR), 1775–1784

Yang, C., Fang, M., Zheng, X. Z., et al. 2026, AJ, 171, 238, doi: 10.3847/1538-3881/ae4b36