



Reducing Ground-based Astrometric Errors with Gaia and Gaussian Processes

W. F. Fortino^{1,2}, G. M. Bernstein¹, P. H. Bernardinelli¹, M. Agüena^{3,4}, S. Allam⁵, J. Annis⁵, D. Bacon⁶, K. Bechtol⁷, S. Bhargava⁸, D. Brooks⁹, D. L. Burke^{10,11}, J. Carretero¹², A. Choi¹³, M. Costanzi^{14,15}, L. N. da Costa^{4,16}, M. E. S. Pereira¹⁷, J. De Vicente¹⁸, S. Desai¹⁹, P. Doel⁹, A. Drlica-Wagner^{5,20,21}, K. Eckert¹, T. F. Eifler^{22,23}, A. E. Evrard^{17,24}, I. Ferrero²⁵, J. Frieman^{5,21}, J. García-Bellido²⁶, E. Gaztañaga^{27,28}, D. W. Gerdes^{17,24}, R. A. Gruendl^{29,30}, J. Gschwend^{4,16}, G. Gutierrez⁵, W. G. Hartley³¹, S. R. Hinton³², D. L. Hollowood³³, K. Honscheid^{14,34}, D. J. James³⁵, M. Jarvis¹, S. Kent^{5,21}, K. Kuehn^{36,37}, N. Kuropatkin⁵, M. A. G. Maia^{4,16}, J. L. Marshall³⁸, F. Menanteau^{29,30}, R. Miquel^{13,39}, R. Morgan⁷, J. Myles⁴⁰, R. L. C. Ogando^{4,16}, A. Palmese^{5,21}, F. Paz-Chinchón^{30,41}, A. A. Plazas⁴², A. Roodman^{10,11}, E. S. Rykoff^{10,11}, E. Sanchez¹⁸, B. Santiago^{4,43}, V. Scarpine⁵, M. Schubnell¹⁷, S. Serrano^{27,28}, I. Sevilla-Noarbe¹⁸, M. Smith⁴⁴, E. Suchyta⁴⁵, G. Tarle¹⁷, C. To^{10,11,40}, D. L. Tucker⁵, T. N. Varga^{46,47}, A. R. Walker⁴⁸, J. Weller^{46,47}, and W. Wester⁵

(The DES Collaboration)

¹ Department of Physics and Astronomy, University of Pennsylvania, Philadelphia, PA 19104, USA; fortino@sas.upenn.edu² Department of Physics and Astronomy, University of Delaware, Newark, DE 19716, USA; garyb@physics.upenn.edu³ Departamento de Física Matemática, Instituto de Física, Universidade de São Paulo, CP 66318, São Paulo, SP, 05314-970, Brazil⁴ Laboratório Interinstitucional de e-Astronomia—LInEA, Rua Gal. José Cristino 77, Rio de Janeiro, RJ—20921-400, Brazil⁵ Fermi National Accelerator Laboratory, P.O. Box 500, Batavia, IL 60510, USA⁶ Institute of Cosmology and Gravitation, University of Portsmouth, Portsmouth, PO1 3FX, UK⁷ Physics Department, 2320 Chamberlin Hall, University of Wisconsin-Madison, 1150 University Avenue, Madison, WI 53706-1390, USA⁸ Department of Physics and Astronomy, Pevensey Building, University of Sussex, Brighton, BN1 9QH, UK⁹ Department of Physics & Astronomy, University College London, Gower Street, London, WC1E 6BT, UK¹⁰ Kavli Institute for Particle Astrophysics & Cosmology, P.O. Box 2450, Stanford University, Stanford, CA 94305, USA¹¹ SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA¹² Institut de Física d'Altes Energies (IFAE), The Barcelona Institute of Science and Technology, Campus UAB, E-08193 Bellaterra (Barcelona), Spain¹³ Center for Cosmology and Astro-Particle Physics, The Ohio State University, Columbus, OH 43210, USA¹⁴ INAF-Osservatorio Astronomico di Trieste, via G.B. Tiepolo 11, I-34143 Trieste, Italy¹⁵ Institute for Fundamental Physics of the Universe, Via Beirut 2, I-34014 Trieste, Italy¹⁶ Observatório Nacional, Rua Gal. José Cristino 77, Rio de Janeiro, RJ—20921-400, Brazil¹⁷ Department of Physics, University of Michigan, Ann Arbor, MI 48109, USA¹⁸ Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Madrid, Spain¹⁹ Department of Physics, IIT Hyderabad, Kandi, Telangana 502285, India²⁰ Department of Astronomy and Astrophysics, University of Chicago, Chicago, IL 60637, USA²¹ Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, USA²² Department of Astronomy/Steward Observatory, University of Arizona, 933 North Cherry Avenue, Tucson, AZ 85721-0065, USA²³ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA²⁴ Department of Astronomy, University of Michigan, Ann Arbor, MI 48109, USA²⁵ Institute of Theoretical Astrophysics, University of Oslo, P.O. Box 1029 Blindern, NO-0315 Oslo, Norway²⁶ Instituto de Física Teórica UAM/CSIC, Universidad Autónoma de Madrid, E-28049 Madrid, Spain²⁷ Institut d'Estudis Espacials de Catalunya (IEEC), 08034 Barcelona, Spain²⁸ Institute of Space Sciences (ICE, CSIC), Campus UAB, Carrer de Can Magrans, s/n, E-08193 Barcelona, Spain²⁹ Department of Astronomy, University of Illinois at Urbana-Champaign, 1002 W. Green Street, Urbana, IL 61801, USA³⁰ National Center for Supercomputing Applications, 1205 West Clark Street, Urbana, IL 61801, USA³¹ Département de Physique Théorique and Center for Astroparticle Physics, Université de Genève, 24 quai Ernest Ansermet, CH-1211 Geneva, Switzerland³² School of Mathematics and Physics, University of Queensland, Brisbane, QLD 4072, Australia³³ Santa Cruz Institute for Particle Physics, Santa Cruz, CA 95064, USA³⁴ Department of Physics, The Ohio State University, Columbus, OH 43210, USA³⁵ Center for Astrophysics | Harvard & Smithsonian, 60 Garden Street, Cambridge, MA 02138, USA³⁶ Australian Astronomical Optics, Macquarie University, North Ryde, NSW 2113, Australia³⁷ Lowell Observatory, 1400 Mars Hill Road, Flagstaff, AZ 86001, USA³⁸ George P. and Cynthia Woods Mitchell Institute for Fundamental Physics and Astronomy, and Department of Physics and Astronomy, Texas A&M University, College Station, TX 77843, USA³⁹ Institució Catalana de Recerca i Estudis Avançats, E-08010 Barcelona, Spain⁴⁰ Department of Physics, Stanford University, 382 Via Pueblo Mall, Stanford, CA 94305, USA⁴¹ Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK⁴² Department of Astrophysical Sciences, Princeton University, Peyton Hall, Princeton, NJ 08544, USA⁴³ Instituto de Física, UFRGS, Caixa Postal 15051, Porto Alegre, RS—91501-970, Brazil⁴⁴ School of Physics and Astronomy, University of Southampton, Southampton, SO17 1BJ, UK⁴⁵ Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831, UK⁴⁶ Max Planck Institute for Extraterrestrial Physics, Giessenbachstrasse, D-85748 Garching, Germany⁴⁷ Universitäts-Sternwarte, Fakultät für Physik, Ludwig-Maximilians Universität München, Scheinerstr. 1, D-81679 München, Germany⁴⁸ Cerro Tololo Inter-American Observatory, NSF's National Optical-Infrared Astronomy Research Laboratory, Casilla 603, La Serena, Chile

Received 2020 October 28; revised 2021 May 28; accepted 2021 May 31; published 2021 August 16

Abstract

Stochastic field distortions caused by atmospheric turbulence are a fundamental limitation to the astrometric accuracy of ground-based imaging. This distortion field is measurable at the locations of stars with accurate positions provided by the Gaia DR2 catalog; we develop the use of Gaussian process regression (GPR) to interpolate the distortion field to arbitrary locations in each exposure. We introduce an extension to standard GPR

techniques that exploits the knowledge that the 2D distortion field is curl-free. Applied to several hundred 90 s exposures from the Dark Energy Survey as a test bed, we find that the GPR correction reduces the variance of the turbulent astrometric distortions $\approx 12\times$, on average, with better performance in denser regions of the Gaia catalog. The rms per-coordinate distortion in the *riz* bands is typically ≈ 7 mas before any correction and ≈ 2 mas after application of the GPR model. The GPR astrometric corrections are validated by the observation that their use reduces, from 10 to 5 mas rms, the residuals to an orbit fit to *riz*-band observations over 5 yr of the $r = 18.5$ trans-Neptunian object Eris. We also propose a GPR method, not yet implemented, for simultaneously estimating the turbulence fields and the 5D stellar solutions in a stack of overlapping exposures, which should yield further turbulence reductions in future deep surveys.

Unified Astronomy Thesaurus concepts: [Astrometry \(80\)](#); [Sky noise \(1463\)](#); [Astronomy data analysis \(1858\)](#)

1. Introduction

Ground-based astrometric measurements are one of the oldest human quantitative scientific endeavors. The accuracy of astrometric data was limited by the resolution of human vision, and subsequently by the angular resolution of telescopes. The successful detection of stellar parallax in the mid-1800s required understanding of stellar aberration and atmospheric refraction, which can be ameliorated to a great extent by differential measurements with respect to nearby more distant stars and attention to observing techniques (e.g., transit telescopes). In the modern era of subarcsecond seeing and digital detectors, the error budget for relative astrometric accuracy of unresolved sources within an exposure will usually be dominated by three contributions:

1. Shot noise: stochastic errors in centroiding of the source in pixel coordinates. This component is typically $\sigma_x \approx \sigma_{\text{PSF}}/\nu$, where σ_{PSF} is the rms width of the point-spread function and ν is the signal-to-noise ratio (S/N) of the detection.⁴⁹
2. Solution accuracy: the static errors in the map from pixel coordinates to (relative) sky coordinates, i.e., the distortions in the optics and detector, and the static refraction of the atmosphere, including chromatic distortions from the atmosphere and optics.
3. Atmospheric turbulence: stochastic wander of the source due to refraction by atmospheric density fluctuations.

The transfer of the relative astrometry of an image to absolute sky coordinates is further limited by the accuracy of the reference catalog used to make such a transfer (as well as the contributions of items 1 and 3 to the exposure's measurements of the reference stars).

While specialized instruments can be designed to improve solution accuracy, we demonstrated in Bernstein et al. (2017, hereafter B17) that an astrometric solution with ≈ 1 mas accuracy is possible for a general-purpose wide-field imager, the Dark Energy Camera (DECam; Flaugher et al. 2015) on the 4 m Blanco Telescope at Cerro Tololo Inter-American Observatory (CTIO). This leaves shot noise (item 1) and atmospheric turbulence (item 3) as the dominant sources of astrometric noise, with the former dominant for faint sources and the latter dominant for high-S/N point sources.

Astrometric science has been revolutionized by space observatories, particularly the Hipparcos (Perryman et al. 1997) and Gaia DR2 (Gaia Collaboration et al. 2018) catalogs. Space telescopes gain substantially in shot-noise errors if diffraction-limited resolution yields low σ_{PSF} ; their stability,

specialized instrumentation, and greatly reduced chromatic effects improve solution accuracy; and, perhaps most importantly, they are free of atmospheric turbulence errors.

The advent of Gaia DR2 also revolutionizes the potential of ground-based astrometry. Most obviously, the density of the DR2 catalog ($O(1)$ star per arcmin²) allows almost any ground-based exposure of modest field of view (FOV) to be placed onto the absolute reference frame of Gaia DR2, obliterating the distinction between absolute and relative astrometry. Typical position/parallax uncertainties in Gaia DR2 rise from ≈ 0.02 mas at $G = 13$ mag to ≈ 2 mas at the catalog limit of $G = 21$ mag (Gaia Collaboration et al. 2018), 1–3 orders of magnitude lower than previous astrometric catalogs approaching similar sky density.

In this paper we demonstrate and quantify another important benefit that Gaia DR2 bestows on ground-based astrometry: the $\approx 1'$ typical spacing between Gaia stars is well below the $\approx 10'$ coherence length of atmospheric turbulence (B17), which means that we can use Gaia as a reference to measure and correct the majority of the power spectrum of astrometric distortions imposed by atmospheric turbulence.

The idea of exploiting the finite correlation length of atmospheric turbulence to reduce the error induced on target stars' positions has been discussed before. This has been of particular interest in astrometric searches for exoplanets (including use with adaptive optics and interferometers). Lazorenko & Lazorenko (2004) propose a fairly complex method to interpolate turbulence to a single target star from an ensemble of nearby reference stars.⁵⁰ This method was applied to exposures from the 8 m Very Large Telescope, yielding estimates of parallax accuracy of 0.04 mas for stars at 17–19 mag (Lazorenko et al. 2009) for exposures accumulating to ≈ 1800 s. (All models and data predict turbulence residuals to decline with exposure time as $T^{-1/2}$.)

We address in this work the application to wide-field surveys, where we are interested in estimating positions for all targets in the field, ideally to the photon-noise limit. The characteristics of atmospheric turbulence were investigated theoretically by Lindegren (1980) and empirically by Han & Gatewood (1995) and Zacharias (1996), among others. B17 found statistics for turbulent distortions to be in rough agreement with these earlier works and proposed the use of Gaussian process regression (GPR) to transfer the turbulent field from Gaia stars to targets of interest. If the turbulence gives rise to a projected (2D) time-delay surface that can be accurately described as a Gaussian random field, then GPR, being the maximum likelihood (ML) estimator for a Gaussian

⁴⁹ Throughout this paper, we will quote astrometric errors or image sizes as the rms per axis on the sky.

⁵⁰ The method of Lazorenko & Lazorenko (2004) may be equivalent to the GPR solution in some limits; we have not investigated this carefully.

process (GP), is also the minimum-variance unbiased interpolator.

Earlier explorations of extraction of precision photometry from wide-field CCD imaging in targeted fields include Platais et al. (2002), Anderson et al. (2006), and Bouy et al. (2013). Typically, any time-dependent distortions such as atmospheric turbulence are corrected through polynomial fits to per-exposure distortions over the span of a single CCD (scales of order $10'$). Anderson et al. (2006) take the additional step of referencing each star to a locally linear transformation determined from ≈ 50 neighboring reference stars, obtaining ≈ 7 mas residuals in their 900 s exposures on the 2.2 m ESO telescope.

While this paper was under review, Lubow et al. (2021) reported application of a similar method—local coordinate systems defined by the 33 nearest Gaia DR2 stars—to the catalogs of the PanSTARRS1 survey, attaining median 2D differentials between PS1 and Gaia positions of ≈ 5 mas with exposure times of 30–45 s on the 1.8 m telescope for stars near $i = 17$ mag. Some comparison to our results is made in Section 7.

In this paper we pursue the application of GPR astrometric interpolation to positional catalogs from exposures in the Dark Energy Survey (DES; Diehl et al. 2014) and develop a method that can be applied in “production mode” to the $O(10^5)$ exposures and $O(10^9)$ detections of unresolved sources in that survey. As reported by B17, the 90 s DES exposures exhibit strongly anisotropic stochastic distortions with typical rms amplitudes of 5–10 mas. This dominates the ≈ 1 mas systematic errors in the calibration of the DECam astrometric map. In this work we will demonstrate that GPR from Gaia DR2 stars succeeds in reducing the rms stochastic distortions to ≈ 2 mas per axis.

This greatly surpasses the requirement set for the Vera C. Rubin Observatory of ≤ 10 mas rms relative astrometric accuracy per axis.⁵¹ GPR turbulence reduction will allow the Rubin Observatory Legacy Survey of Space and Time (LSST)⁵² to push astrometric science further beyond the capabilities of Gaia in many ways. The LSST will be able to measure Gaia-quality stellar parallaxes/proper motions well beyond Gaia’s faint limit, as well as improving on Gaia accuracy for stars near its limit. LSST can also bring milliarcsecond precision to the tracking of minor planets and other transients.

In the next section we review the relevant aspects of DES imaging and astrometric results from B17. Section 3 reviews standard GPRs. Astrometric interpolation differs from standard cases in that the turbulent image displacement (u, v) is observed to follow the expectation that it is curl-free. We show how to extend the GPR formalism to exploit this known relation between the u and v fields.

Section 4 derives the correlation function—or, in GPR parlance, the “kernel”—that should result from wind-blown von Karman turbulence at a single layer of the atmosphere. Section 5 describes the numerical methods for choosing and applying a kernel to the catalogs, and Section 6 gives quantitative results of application of the curl-free GPR to a test sample of several hundred DES exposures, including validation by fitting an orbit to DES observations of the bright trans-Neptunian object Eris. We conclude in Section 7.

⁵¹ See Table 18 of <https://docushare.lsstcorp.org/docushare/dsweb/Get/LPM-17>.

⁵² <https://www.lsst.org>

In the Appendix, we derive an even more comprehensive use of the GPR methodology, in which one can simultaneously obtain the ML values for the distortion fields of a stack of exposures and the 5D position/parallax/proper-motion solutions of the stars contained within this stack. This method has the potential for significant further reduction of turbulence residuals, by effectively turning every star in the field with (shot noise) \lesssim (turbulence noise) into a reference star, not just those with Gaia measures. For DES exposures, the density of such stars does not exceed the density of Gaia stars, so the benefit of this technique will not be large, and we have not implemented this method on DES data. It should be of more value to future surveys with larger telescopes.

Table 1 is a guide to the notation in this paper.

2. Summary of DES Data and Astrometry

2.1. DES Data and Astrometric Solution

The DECam science array consists of 62 distinct CCDs, each 2048×4096 pixels at $\approx 0''.263$ per pixel. The FOV approximates a circle with $\approx 1^\circ$ radius. The analyses in this paper are done on exposures taken as part of the “Wide” survey of DES, in which a 5000 deg^2 section of the southern galactic cap is imaged 10 times in each of the $g, r, i, z,$ and Y filter bands, spread over six annual August–February observing seasons. We will for the most part ignore the Y -band exposures, which have substantially lower S/N than $griz$ and will not contribute significantly to overall astrometric precision (but appear otherwise astrometrically well-behaved). The $griz$ exposures for the Wide survey are all 90 s in duration.

The results reported here make use of the “Y6A1” internal release of the full Wide survey data. The individual exposures are processed with the “FinalCut” pipeline very similar to the earlier version described in Morganson et al. (2018). Pixel coordinates for all sources are determined from the [XY] WIN_IMAGE windowed centroid quantity measured by SEExtractor (Bertin & Arnouts 1996). We assign the ERRA-WIN_IMAGE measurement as the σ of a circular Gaussian measurement error on each unresolved source.

A mapping from pixel coordinates to sky coordinates is derived using the methods described in B17, with some improvements. The astrometry solution includes these terms:

1. A cubic polynomial spanning the whole FOV for each exposure, which absorbs the static atmospheric refraction, stellar aberration, and a pointing solution.
2. A zenith-oriented differential chromatic refraction term.
3. A polynomial per CCD per band per observing season, which captures optical distortions.
4. A chromatic lateral color shift oriented radially.
5. * Short-timescale (weeks) affine shifts in the positions of the CCDs in the focal plane.
6. The “tree ring” and “glowing edge” distortions arising from stray electric fields in each CCD.
7. * An additional fixed map of distortions apparently due to electric fields around the electrical cable connector.

Those items marked with an asterisk have been added to the model since B17. Furthermore, the entire astrometric solution is now registered to Gaia DR2 and allows for nonzero proper motion and parallax for all stars in the DES footprint when registering images internally and to Gaia. In B17 it was demonstrated that any errors in the astrometric solution that repeat over time are limited to $\lesssim 1$ mas rms.

Table 1
Notation Used Throughout the Text

Symbol	Meaning
$A(\mathbf{x})$	Aperture function of the telescope
D	Telescope diameter
$d = D/h$	Angular size of the telescope diameter as projected at the turbulent layer
h	Height of turbulence
J	Bessel function of the first kind
\mathbf{K}	Covariance matrix of the GP
$K_{ij} = K(\mathbf{x}_i - \mathbf{x}_j; \boldsymbol{\pi}_K)$	Kernel function of the GP
$K_r(\Delta\mathbf{x})$	2×2 covariance function for stars i, j
$K_\phi(\Delta\mathbf{x})$	Two-point correlation function of ϕ
$\mathbf{k} = (k_x, k_y)$	Fourier inverse to $\Delta\mathbf{x}$
k_0	$2\pi h/r_0$
$(k_{\parallel}, k_{\perp})$	Components of \mathbf{k} relative to parallactic direction of the turbulence
N_s	Number of stars in a DES exposure
$\mathbf{n} = \{n_1, n_2, \dots\}$	Set of measurement shot noise for an exposure
$P_\phi(\mathbf{k})$	Power spectrum of ϕ
r_0	Outer scale of turbulence
T	Exposure time
\mathbf{U}	Merged vector of length $2N_s$ of both u and v points
$\mathbf{u} = (u, v)$	Displacements of stellar positions due to atmospheric turbulence
$\tilde{\mathbf{u}}_*$	Posterior predictive mean of the GP regressor
$W(\mathbf{x})$	Line segment with length and direction wT
\mathbf{w}	Wind vector at the turbulence layer
$\mathbf{x} = (x, y)$	Position of stars on a local projection of the sky
$\Delta\mathbf{x}$	$\mathbf{x}_i - \mathbf{x}_j$ for some stars i, j
z	Zenith angle of the telescope
$\boldsymbol{\mu}$	Mean of a GP, taken to be 0
ν	S/N of a source detection
ξ_0	Overall amplitude of the turbulence = $(\xi_{uu} + \xi_{vv})(\mathbf{x} = 0)$
$\xi_{uu}, \xi_{uv}, \xi_{vv}$	Two-point correlation functions among the (u, v) components of displacement
ξ_+	$\xi_{uu} + \xi_{vv}$
$\boldsymbol{\pi}_K$	Hyperparameters for the kernel function
σ	Stochastic errors in the centroiding of a source in pixel coordinates
σ_{PSF}	rms width of the point-spread function
$\phi(\mathbf{x})$	OPD along the line of sight

2.2. Nature of the Stochastic Distortions

Aside from the per-exposure cubic polynomial spanning the full FOV, these baseline Y6A1 astrometric solutions do not attempt to remove any of the stochastic distortions that would arise from atmospheric turbulence or other effects varying on timescales of single exposures. B17 described several properties of the stochastic distortion patterns, which we summarize here.

To quantify the distortion fields, we first define (u, v) to be the components of the difference between the DES-derived R.A./decl. and the Gaia DR2 values. We will summarize the distortion field with the two-point correlation function of this error field, defined as

$$\xi_{uu}(\mathbf{x}) = \langle u_i u_j \rangle, \quad (1)$$

where i and j range over all stars separated by the vector \mathbf{x} . The virtue of this statistic is that the contribution from shot noise—or any other form of noise that has negligible star-to-star correlation—averages to zero. We can similarly define ξ_{vv} , ξ_{uv} . Of particular interest is

$$\xi_+(\mathbf{x}) \equiv \xi_{uu}(\mathbf{x}) + \xi_{vv}(\mathbf{x}). \quad (2)$$

As $\mathbf{x} \rightarrow 0$, the value of ξ_{uu} yields the total rms variance in the u direction caused by atmospheric turbulence (or other spatially correlated errors).

The DES exposures are seen to have anisotropic errors, i.e., $\xi_{uu} \neq \xi_{vv}$ and $\xi_{uv} \neq 0$. Figure 1 shows the distribution of major/minor axes of the error ellipse implied by the zero-lag limits of the ξ values. The modal major and minor axes are 7 and 5 mas, respectively, with the means being higher. This is the turbulence noise.

A further critical observation is that the stochastic distortion field is curl-free, as seen in Figure 2. This implies that the scalar $\xi_+(\mathbf{x})$ function is a complete description of the vector turbulent distortions, if they arise from a GP, as explained in the Appendix of BK17. As will be detailed below, this function indicates a typical correlation length of $\approx 6'$ for the stochastic distortions (where ξ_+ drops to half its $\mathbf{x} = 0$ value).

3. Curl-free, Anisotropic Gaussian Process Interpolation

3.1. Scalar Gaussian Process Review

The standard methodology for GPR follows from the assumption that we are interested in the value of some stochastic scalar $u(\mathbf{x})$ over the field \mathbf{x} (in our case, the positions $\mathbf{x} = (x, y)$ of a star on a local projection of the sky). The generation of u is considered a GP if and only if the distribution of u at any collection of points $\mathbf{x} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$ can be




Figure 1. Distribution of the major and minor axes of the turbulence-induced astrometric errors for the $\approx 80,000$ exposures of the DES Y6A1 Wide survey. The distributions are clipped at 3 mas since smaller values are unreliable in these data.




Figure 2. The divergence and curl of the astrometric residuals on exposure #228645 (30 s, z band) are plotted on a common scale. The continuity of the vector field across chip boundaries, the curl-free nature of the field, and the streaky pattern of divergence suggest the hypothesis that these distortions arise from wind-blown atmospheric turbulence. Reproduced from BK17.

described by a multivariate normal distribution⁵³

$$\mathbf{u} \equiv \{u_1, u_2, \dots, u_N\} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{K}), \quad (3)$$

with a mean $\boldsymbol{\mu}$ that we can usually take as zero and a covariance matrix $\mathbf{K} \equiv \langle \mathbf{u}\mathbf{u}^T \rangle$ with a known form $K_{ij} = K(\mathbf{x}_i, \mathbf{x}_j)$. We will consider stationary (but potentially anisotropic) fields with possible hyperparameters $\boldsymbol{\pi}_K$, for which we can place $K_{ij} = K(\mathbf{x}_i - \mathbf{x}_j; \boldsymbol{\pi}_K)$.

If we consider the field to have known values \mathbf{u} at training points \mathbf{x} , and we seek estimates \mathbf{u}^* at points \mathbf{x}^* , then we can write the joint distribution as

$$\begin{pmatrix} \mathbf{u} \\ \mathbf{u}^* \end{pmatrix} \sim \mathcal{N}\left(0, \begin{bmatrix} \mathbf{K} & \mathbf{K}^* \\ \mathbf{K}^{*T} & \mathbf{K}^{**} \end{bmatrix}\right). \quad (4)$$

Here we follow GP convention by setting $K_{ij}^* = K(\mathbf{x}_i - \mathbf{x}_j^*)$ and $K_{ij}^{**} = K(\mathbf{x}_i^* - \mathbf{x}_j^*)$. Conditioning this joint distribution on the measured \mathbf{u} yields the standard scalar GPR solution:

$$\mathbf{u}^* | \mathbf{u} \sim \mathcal{N}(\mathbf{K}^{*T} \mathbf{K}^{-1} \mathbf{u}, \mathbf{K}^{**} - \mathbf{K}^{*T} \mathbf{K}^{-1} \mathbf{K}^*). \quad (5)$$

Thus, the ML solution is the mean of this conditioned distribution—the ML value of \mathbf{u}^* at any point is a linear combination of the values at \mathbf{u} , with coefficients determined by the relative locations of the reference stars to the target

position. The presence of independent measurement noise $\mathbf{n} = \{n_1, n_2, \dots\}$ on the training points can be incorporated by considering the noise to be an addition to the kernel/covariance matrix \mathbf{K} that has zero correlation length, i.e., only appears on the diagonal

$$K_{ij} \rightarrow K(\mathbf{x}_i - \mathbf{x}_j) + \delta_{ij} \langle n_i^2 \rangle. \quad (6)$$

The suitability of the chosen kernel to the training data is quantified by the log marginal likelihood,

$$-2 \log p(\mathbf{u} | \mathbf{x}, \boldsymbol{\pi}_K) = \mathbf{u}^T \mathbf{K}^{-1} \mathbf{u} + \log |2\pi \mathbf{K}|, \quad (7)$$

where \mathbf{K} is a function of the hyperparameters $\boldsymbol{\pi}_K$. The most suitable kernel is chosen by maximizing this quantity over $\boldsymbol{\pi}_K$. If the field is truly Gaussian, the resultant interpolator is then optimal. The GPR does, however, yield a functional interpolator even if \mathbf{K} is not precisely the covariance function of \mathbf{u} , or if \mathbf{u} is not a Gaussian field.

Application of the standard GPR to ground-based astrometric data proceeds by using the sky positions of the stars as the feature vectors \mathbf{x}_i and defining two distinct scalar fields (u, v) as the differences between the observed x and y coordinates and the true coordinates, $(u_i, v_i) = \mathbf{x}_i^{\text{obs}} - \mathbf{x}_i^{\text{true}}$. The training data are those for which the Gaia DR2 motions are available (and calculated at the observation epoch). Distinct GPRs are trained for each component u and v in this simplest scheme.

3.2. Curl-free Vector Fields

A simple and applicable ray-optic model for the displacements caused by atmospheric turbulence is that the centroid of the stellar image moves by an apparent vector

$$\mathbf{u} = (u, v) = \nabla_x \phi(\mathbf{x}), \quad (8)$$

where $\phi(\mathbf{x})$ is the optical path difference (OPD) along the line of sight \mathbf{x} , as convolved with the telescope aperture. If this model holds instantaneously and the source photon arrival rate is constant, then we can average both \mathbf{u} and ϕ over the duration of the exposure and the equation will still hold. A special case would be a single-screen “frozen” approximation: $\phi(\mathbf{x}, t)$ varies with time as $\phi(\mathbf{x} - \mathbf{w}t)$ for a wind vector \mathbf{w} at the turbulent layer. In this case, the time-averaged ϕ is the convolution of the instantaneous ϕ with a line segment of wind motion during the exposure. The results in this section will, however, be true for any model of the time-averaged ϕ .

As shown in B17, the two-point correlation functions of the residual displacement fields $\mathbf{u}(\mathbf{x})$ are observed to be purely “E-mode,” meaning that $\nabla \times \mathbf{u}$ contains only white noise consistent with the shot noise of stellar centroid measurements. It is therefore likely that enforcing a curl-free turbulence field will yield a more accurate GPR, since we can combine the data from both fields u and v to solve for a single degree of freedom ϕ .

The curl-free GPR begins by assuming that ϕ is a zero-mean Gaussian field, in which case it can be fully characterized by its power spectrum $P_\phi(\mathbf{k})$ or its two-point correlation function $K_\phi(\Delta\mathbf{x})$, which we can relate via the Fourier transform

$$K_\phi(\Delta\mathbf{x}) = \int d^2\mathbf{k} P_\phi(\mathbf{k}) e^{i\mathbf{k} \cdot \Delta\mathbf{x}}. \quad (9)$$

The relation (8) implies that the Fourier transforms of \mathbf{u} and ϕ satisfy $FT(\mathbf{u}) = i\mathbf{k} FT(\phi)$. This in turn implies that the

⁵³ Using the standard notation $\mathcal{N}(\boldsymbol{\mu}, C)$ for a multivariate normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix C .

covariance matrix of the turbulence component of the displacement field is

$$K_t(\Delta\mathbf{x}) = \int d^2\mathbf{k}(\mathbf{k}\mathbf{k}^T)P_\phi(\mathbf{k})e^{i\mathbf{k}\cdot\Delta\mathbf{x}}. \quad (10)$$

Thus, any parametric model for P_ϕ can be transformed into a parametric 2×2 covariance matrix function of $\Delta\mathbf{x}$. (We retain a scalar notation for this 2×2 function K_t , reserving the matrix symbol \mathbf{K} for the matrix of covariances between different field points.)

With this modification, we can recast the GPR to a merged vector $\mathbf{U} = \{u_1, v_1, u_2, v_2, \dots, u_{N_s}, v_{N_s}\}$ of $2N_s$ scalar quantities. As in the scalar GPR, we can define matrices \mathbf{K} , \mathbf{K}^* , \mathbf{K}^{**} composed of 2×2 blocks such that

$$\begin{aligned} K_{i,j} &= K_t(\mathbf{x}_i - \mathbf{x}_j) + \delta_{ij}\langle \mathbf{n}_i \mathbf{n}_j^T \rangle \\ K_{i,j}^* &= K_t(\mathbf{x}_i - \mathbf{x}_j^*) \\ K_{i,j}^{**} &= K_t(\mathbf{x}_i^* - \mathbf{x}_j^*). \end{aligned} \quad (11)$$

The measurement noise term in the first row is also a 2×2 matrix, which in practice is very close to diagonal for DES data but is substantially anisotropic for Gaia. After this change, the standard GPR formulae (5) and (7) remain correct, with the substitution $\mathbf{u} \rightarrow \mathbf{U}$. If the $K_t(\Delta\mathbf{x})$ function were diagonal, then the GPR solution would separate into two distinct scalar GPRs. But even an isotropic field ϕ has off-diagonal terms in K_t , and therefore this curl-free interpolant exploits information that two distinct scalar GPRs cannot.

4. Model for Wind-blown Von Karman Turbulence

We desire a model for the correlation function $K(\mathbf{x})$ between the turbulent distortions at two locations separated by \mathbf{x} . Since the displacement (u,v) is 2D, there are three scalar correlation functions of interest, K_{uu} , K_{uv} , K_{vv} . Considering first one dimension of the deflection, u , we have from Equation (8) that the deflection of a single photon passing through a time-delay screen ϕ at location \mathbf{x} and time t will be $u \propto \frac{\partial}{\partial x}\phi(\mathbf{x}, t)$. The photons that arrive to form the image of a particular star arrive at fixed t from a range of \mathbf{x} described by the aperture function of the telescope $A(\mathbf{x})$, effectively convolving the instantaneous OPD screen as $\phi \rightarrow \phi \otimes A$. Adopting the frozen-screen approximation common in adaptive optics analyses, such that $\phi(\mathbf{x}, t) = \phi(\mathbf{x} - \mathbf{w}t)$, means that the integration of photon arrival time over the exposure duration T corresponds to a further convolution of the phase screen by a line segment of length and direction $\mathbf{w}T$, which we will denote by the ‘‘wind function’’ $W(\mathbf{x})$. The apparent deflection, averaged over all arriving photons for a star crossing the phase screen at location \mathbf{x} , becomes

$$u(\mathbf{x}) \propto \frac{\partial}{\partial x}[\phi(t=0) \otimes A \otimes W]. \quad (12)$$

The convolutions become multiplications in the Fourier domain, and we can also recall that the correlation function K_{uu} is the Fourier transform of the power spectrum $P_u(\mathbf{k})$ of the

deflections, giving

$$\begin{pmatrix} K_{uu}(\mathbf{x}) \\ K_{uv}(\mathbf{x}) \\ K_{vv}(\mathbf{x}) \end{pmatrix} \propto \int d^2\mathbf{k} \begin{pmatrix} k_x^2 \\ k_x k_y \\ k_y^2 \end{pmatrix} P_\phi(\mathbf{k}) |\tilde{A}^2(\mathbf{k})| |\tilde{W}^2(\mathbf{k})| e^{i\mathbf{k}\cdot\mathbf{x}}, \quad (13)$$

which reproduces Equation (10) by the inclusion of the aperture and wind functions altering the instantaneous OPD screen ϕ . In this section we have treated \mathbf{x} as being in distance units and \mathbf{k} in inverse distance. But we can equally well use these equations with \mathbf{x} representing the 2D angle subtended about the telescope axis, which we shall do henceforth.

It now remains to choose models for P_ϕ , A , and W . The common model for the power spectrum of atmospheric turbulence is the von Karman spectrum,

$$P_\phi(k) \propto (k^2 + k_0^2)^{-11/6}, \quad (14)$$

where k_0 is the inverse of the outer scale of turbulence. More precisely, our angular system, with turbulence at height h and outer scale r_0 , produces $k_0 = 2\pi h/r_0$. There is an additional complication that when the zenith angle z of the observation is nonzero, the (horizontal) turbulent layer is foreshortened along the direction toward zenith. To include this effect, we should break \mathbf{k} into components $(k_{\parallel}, k_{\perp})$ relative to the parallactic direction and substitute

$$k^2 \rightarrow (k_{\parallel} \cos z)^2 + k_{\perp}^2 \quad (15)$$

in the argument of P_ϕ .

For the telescope aperture we will adopt the simplest case of a uniformly filled circular aperture, which yields

$$|\tilde{A}^2(\mathbf{k})| \propto \left[\frac{J_1(kd/2)}{kd} \right]^2, \quad (16)$$

where J is the Bessel function of the first kind and d is the angular size of the telescope diameter D as projected to the turbulent layer, $d = D/h$.

The Fourier transform of the line segment window function yields

$$|\tilde{W}^2(\mathbf{k})| \propto \text{sinc}^2(\mathbf{k} \cdot \mathbf{w}/2), \quad (17)$$

where \mathbf{w} is now taken to be the wind transport over the duration of the exposure, as projected onto the angular coordinates of the telescope.

Combining the previous five equations yields our model for the astrometric correlation function $\mathbf{K}(\mathbf{x})$ for a single layer of von Karman turbulence. There are two known parameters (the zenith angle z and the parallactic angle defining k_{\parallel}) and five free parameters: $\{\xi_0, r_0, d, w_x, w_y\}$, denoting the overall amplitude of the turbulence, $\xi_0 = (K_{uu} + K_{vv})(\mathbf{x} = 0)$, the (angular equivalents of) outer scale, telescope aperture, and components of the wind vector. Of these parameters, the outer scale r_0 is the least important because it is typically larger than the 4 m telescope diameter, and the aperture function damps the power spectrum before the outer scale sets in.

The real-space 2D correlation functions for the von Karman spectrum, the aperture, and the wind function can be done analytically individually, but their convolution is not analytic. One has the option of doing the convolution numerically, but we opt instead to use the above equations to multiply the three (analytic) Fourier domain functions and do the Fourier

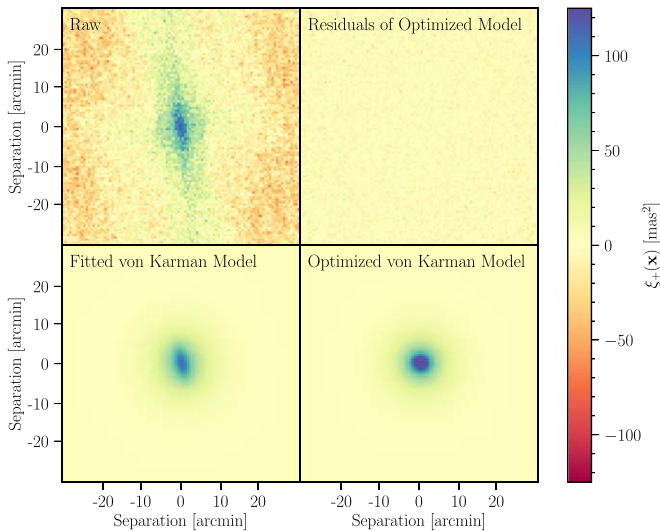


Figure 3. Each panel plots a measure or model of the astrometric error correlation function $\xi_+(\mathbf{x})$. The top left panel shows the measured correlation function for exposure 361582 (in i band), before any GPR subtraction. The bottom left panel shows the single-screen wind-blown von Karman turbulence model that is the best fit to these data for $|\mathbf{x}| < 5'$. The bottom right panel shows the von Karman model that the optimizer finds to minimize the ξ_0 after GPR subtraction. The top right panel shows the measured ξ_+ after subtraction of the GPR. Note that there is no noticeable remaining correlation at any \mathbf{x} (at this dynamic range) after the GPR correction is applied.

transform numerically using a fast Fourier transform. The bottom panels of Figure 3 plot $\xi_+ = K_{uu} + K_{vv}$, for two variants of the wind-blown von Karman model.

We note at this point that the real atmospheric turbulence is unlikely to arise from a single von Karman screen. The observed ξ_+ functions have more structure than the von Karman model, e.g., features at multiple position angles. The turbulence may arise at multiple heights with varying wind speed, for example. We will, however, proceed with this model as a viable kernel, and remember that it is possible that other kernels would perform better, and the most successful von Karman parameters may not represent the physical conditions of a single screen. Data from instruments monitoring atmospheric turbulence are available at CTIO, but we do not expect ab initio information on the turbulence parameters to provide a more effective GP kernel than fitting and optimization to the empirical information from Gaia stars' residuals.

5. Numerical Methods

The curl-free GPR is implemented on the DES data on an exposure-by-exposure basis, with the following procedure:

1. **Data Retrieval:** A Python program retrieves the Y6A1 DES single-epoch astrometric solutions for a single exposure from DES data files, including estimates of shot noise n_i from the SExtractor quantity (ERRAWIN_WORLD). Corresponding Gaia five-parameter astrometry solutions, as well as their full covariance matrices, are retrieved via the Python package ASTROQUERY.
2. **Pre-processing:** The retrieved Gaia 5D solution for epoch J2015.5 for each star is reduced to a 2D position (and uncertainty) at the time of the DES exposure. The DES and Gaia catalogs are then matched with the `match_coordinates_sky` routine from ASTROPY; detections within $0''.5$ are matched. Gaia detections with

no DES match are discarded. Additionally, if more than 15,000 matches are found, 15,000 detections are randomly chosen and the rest are discarded in order to conserve memory in later steps. The program then performs a gnomonic projection on both catalogs using the center of the DES exposure as the projection axis. For each matched star, the residual field $\mathbf{u} = \mathbf{x}_{DES} - \mathbf{x}_{Gaia}$ is formed, and a 2×2 shot-noise uncertainty \mathbf{n}_i is created from the sum of the DES and Gaia measurement errors. In order to facilitate fivefold cross-validation later on, the matched data are randomly distributed into five subsets, A through E. At any one time, one subset will be the validation set, and the other four will be the training set. Additionally, the detections in the DES catalog that did not have a match are designated as the prediction set.

3. **Error Rescaling:** In order to remove outliers, sigma-clipping is performed to four standard deviations on the residual values, \mathbf{u}_i . Additionally, any detections that have shot-noise errors $>250 \text{ mas}^2$ are removed since they are too noisy to be useful and/or are measurement artifacts. For r -, i -, and z -band exposures there are few detections with shot-noise errors this high, but this is an important step for the g and Y bands. We also find that the shot-noise errors from SExtractor are systematically underestimated, particularly for detections with very low shot noise. To give the kernel more accurate shot-noise values, we replace the (ERRAWIN_WORLD) with the rms of the residual field values for a group of 256 detections with similar estimated shot noise. The best-fit third-degree polynomial as a function of astrometric position in the FOV is subtracted from the residual field. This is done to remove large-scale systematic errors from instrument distortion (e.g., thermal expansion of the telescope) and low-altitude “ground-layer” turbulence that distorts large angular scales (see Section 6.2 of B17 for further details).
4. **Fitting kernel to ξ_+ :** In order to get an initial guess of the kernel parameters, π_K , $\xi_+(\mathbf{x})$ is calculated for the training set. A Nelder–Mead optimizer varies π_K to minimize the residual sum of squares (RSS) between the observed ξ_+ and the parametric von Karman $\xi_+ = (K_{uu} + K_{vv})$. For this least-squares fitting, only $\xi_+(\mathbf{x})$ for $x < 5'$ is used. With the resultant kernel parameters, we execute the curl-free GPR on the \mathbf{u}_i . Using fivefold cross-validation, we obtain a residual $\tilde{\mathbf{u}}_i$ for each matched star from the difference of the raw \mathbf{u}_i and a GPR-estimated value that has *not* been trained on that star. Another round of sigma-clipping is then performed on the data, this time removing all detections greater than four standard deviations from the mean of $\tilde{\mathbf{u}}_*$.
5. **Optimization of kernel:** Ideally the GPR kernel fitted to the observed ξ_+ function would be the optimal interpolator, yielding the smallest rms errors in position for the validation set. The rms errors can be measured without shot-noise biases as the limit of $\xi_+(\mathbf{x})$ as $\mathbf{x} \rightarrow 0$. We define a figure of merit meant to approximate this limit as

$$\xi_{1,2} \equiv \langle \xi_+(\mathbf{x}) \rangle_{|\mathbf{x}| < 1/2}, \quad (18)$$

where the average is pair weighted. In practice, the fitted kernel is not optimal for $\xi_{1,2}$, perhaps because the von Karman kernel model does not fully describe the field, or

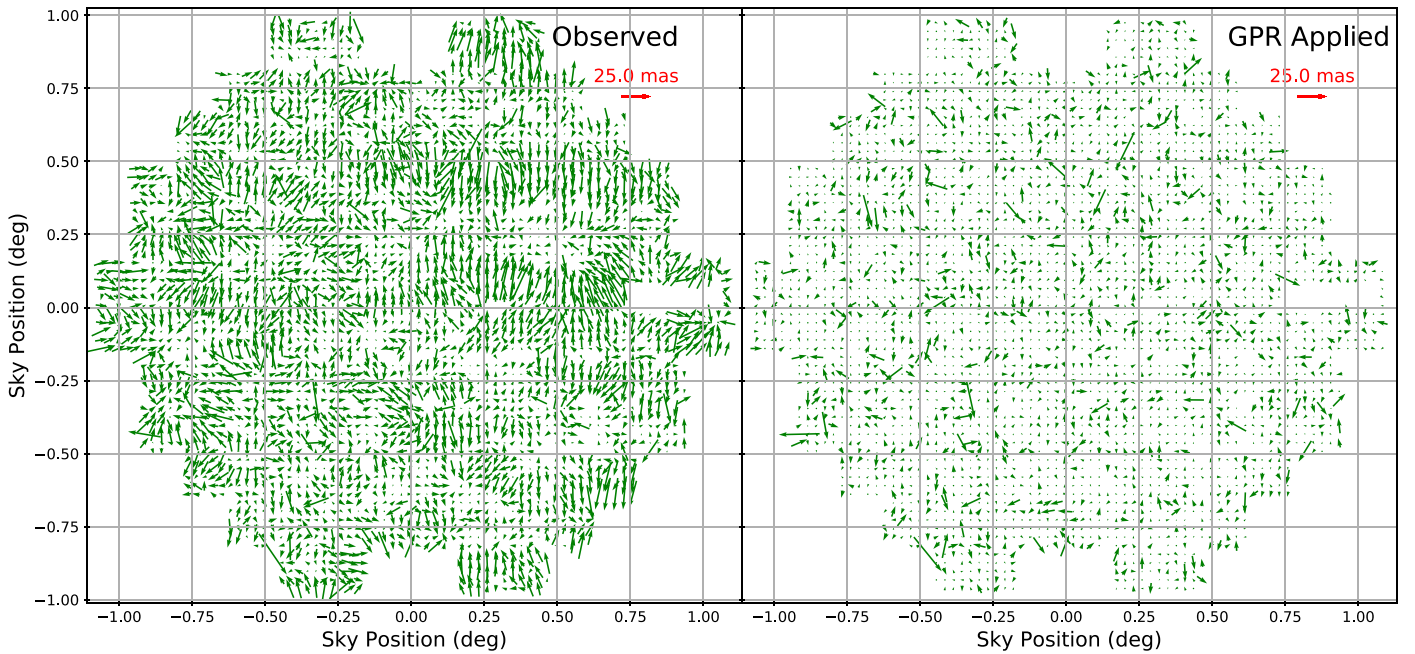


Figure 4. At left is a map of binned astrometric errors of stars in a typical DES exposure (#361582, i band) relative to their Gaia DR2 positions. At right are the results after subtraction of corrections derived from the curl-free GPR method. Fivefold cross-validation is used such that no Gaia star is used to derive the GPR that is applied to it. The reduction in amplitude and correlation length of the astrometric errors is apparent.

perhaps because of non-Gaussianity and/or outliers in the data. Regardless of the cause, we find that the figure of merit $\xi_{1,2}$ can be reduced by further optimization of the von Karman parameters π_K . Using the “fitted” kernel parameters found in the previous step as an initial guess, we use the `L_BFGS_B` gradient-descent method of the `scipy.optimize.minimize` optimizer (Virtanen et al. 2020) to minimize the $\xi_{1,2}$ of the validation set. This is the most computationally intensive part of the process since the optimization requires many repeated evaluations of the GPR. In order to reduce the number of optimization steps, we first optimize over the parameters d , w_x , and w_y and then fix these and optimize the remaining parameters ξ_0 and r_0 . The results π_K we call the “optimized” kernel parameters. We evaluate the GPR and residuals using the optimized parameters. Again using fivefold cross-validation and four standard deviation sigma-clipping, we find $\tilde{\mathbf{u}}_*$ for all matched stars. When we refer to “GPR-subtracted” data, it will by default mean those corrected with the optimized parameters. In some of the further analyses we will explicitly compare results with a GPR using the fitted parameters from step 4 versus the optimized parameters from this step.

6. **Analysis:** With $\tilde{\mathbf{u}}_*$, we calculate statistics on how well the model performs. All statistics omit the stars that have been sigma-clipped.

6. Results for DES Exposures

We applied this procedure to 343 DES exposures (76, 70, 64, 72, and 61 from $griz$ Y bands, respectively). The g -band exposures exhibit very high outlier rates and `ERRWIN_WORLD` inaccuracies for reasons that are not understood and probably irrelevant to the turbulence estimation, so we will not make further use of them in this paper. The low S/N of the Y -

band exposures makes them less interesting tests of turbulence reduction, so most of our statistics will concentrate on riz . Maps of the astrometric errors of Gaia stars before and after application of the GPR correction are shown for a typical exposure in Figure 4.

The average density of matches with Gaia (after sigma-clipping) for the exposures analyzed is 0.95 arcmin^{-2} , or 10^4 per 3 deg^2 DES exposure. With the fivefold cross-validation, there are an average of ≈ 2000 stars in a validation subset. The optimization of the kernel parameters using the `L_BFGS_B` algorithm requires ≈ 100 recalculations of the GPR on average, which is the computational bottleneck of the procedure, requiring several hours on a 12-core cluster node. Further investigation of the optimization may yield substantial speedups.

Our primary measure of success for this analysis is the reduction of the $\xi_0 = \langle u^2 + v^2 \rangle$ of the raw residual field after the GPR model of the turbulence is subtracted. With a finite number of validation stars, we cannot calculate $\xi_0(\mathbf{x} = 0)$. To measure the improvements due to GPR subtraction, we approximate ξ_0 with $\xi_+(\mathbf{x} < 0.5')$, i.e., we average over all validation–star pairs with separation $< 0.5'$.⁵⁴ Figure 5 plots the raw ξ_0 versus the GPR-subtracted ξ_0 for riz Y exposures, with the bottom panel replotting the quantity $\sqrt{\xi_0}/2$, which gives the rms astrometric error per axis. The bottom panel also shows the relationship between rms and field density for riz Y exposures. The average ξ_0 of the raw residual field for all 206 riz -band exposures was 125 mas^2 , or an rms astrometric error of 7.3 mas per axis. The average ξ_0 after the GPR model was subtracted from the residual field was 11.9 mas^2 or 2.3 mas rms. The average reduction in ξ_0 was a factor of 12. This is the principal result of this work.

⁵⁴ Note that this ξ_0 estimator uses a smaller radius ($0.5'$) than the $\xi_{1,2}$ used during kernel optimization. This is because Figure 6 shows that $\xi(\mathbf{x})$ is already dropping significantly at $|\mathbf{x}| = 1.2$ in the GPR-subtracted data.

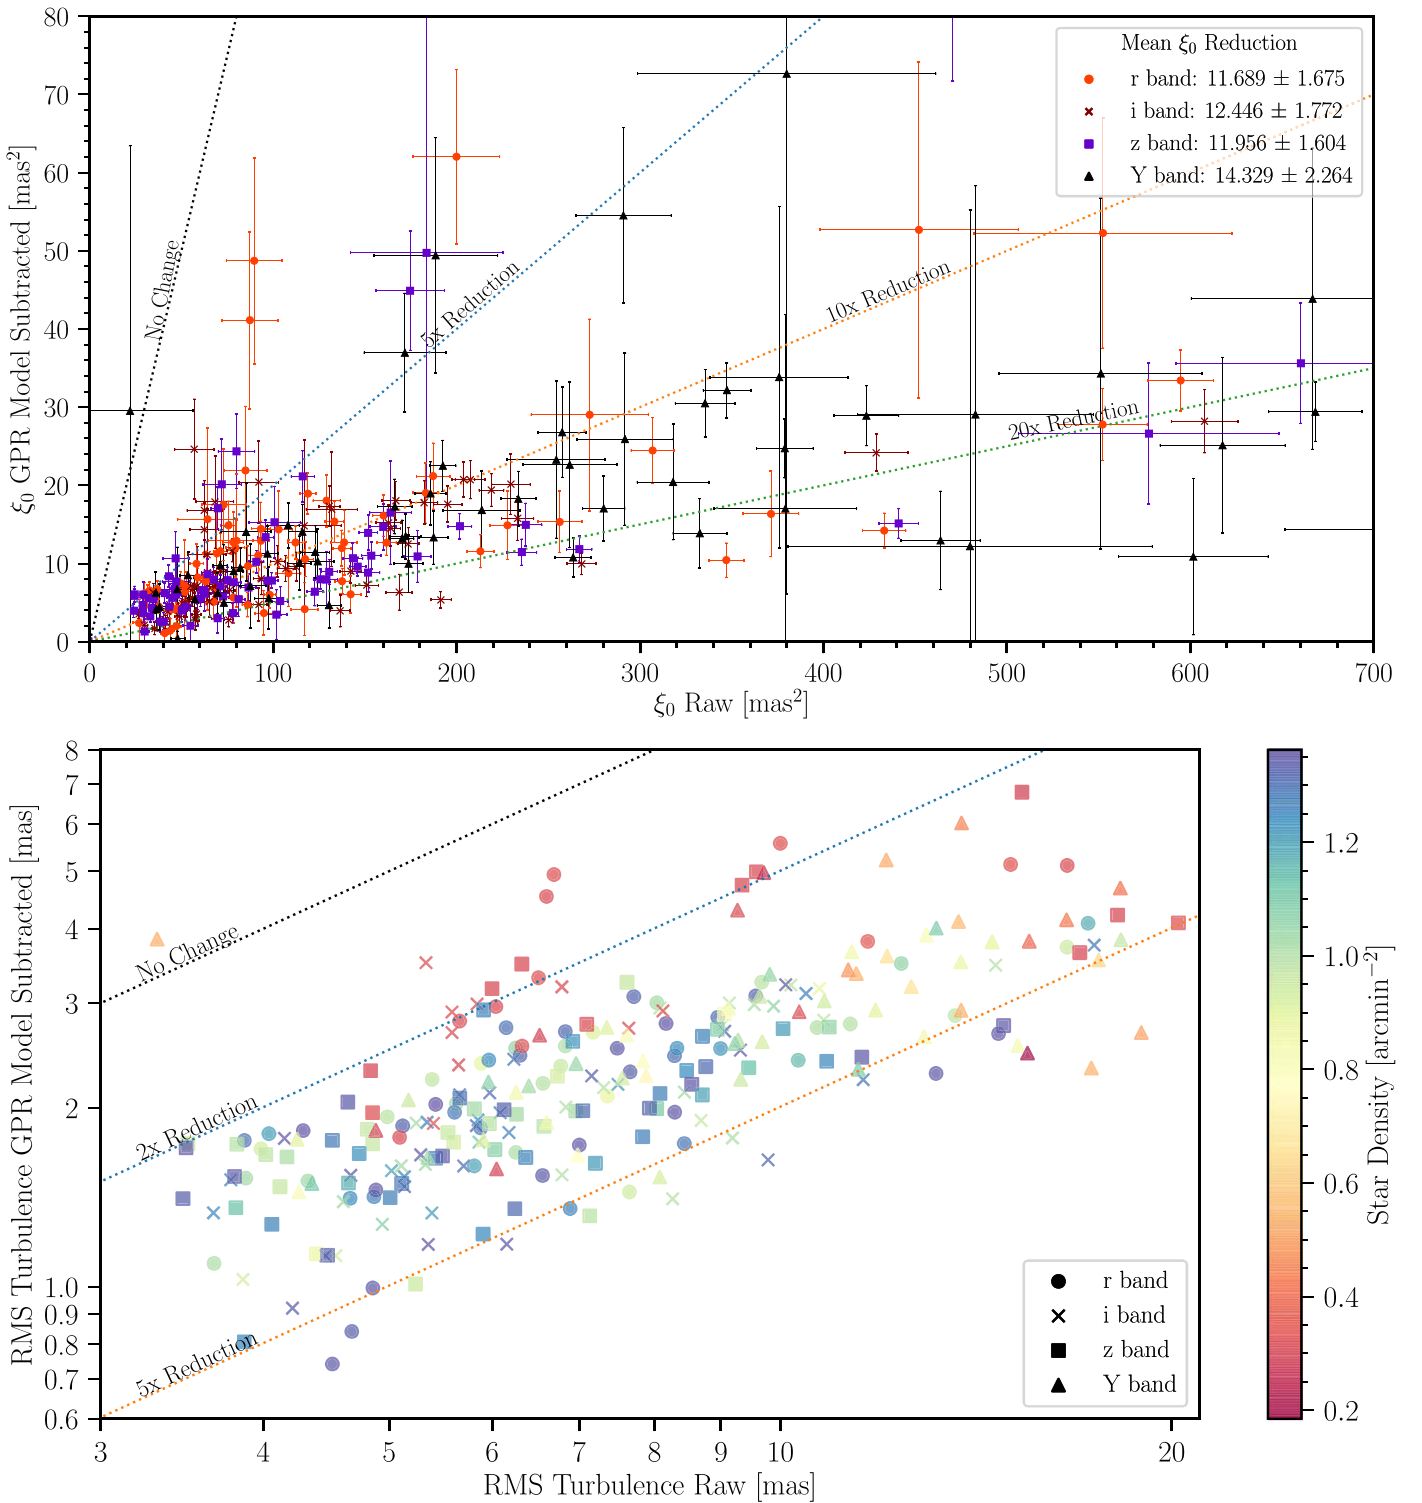


Figure 5. The top panel plots the values of $\xi_0 \approx \xi_+(\lvert\mathbf{x}\rvert < 0.5)$ for the validation samples of every analyzed *riz* *Y* exposure before (*x*-axis) vs. after (*y*-axis) subtraction of the GPR. The dotted lines mark constant factors of ξ_0 reduction by the GPR subtraction. The bottom panel is similar except that it uses $\sqrt{\xi_0}/2$, the rms turbulence error per axis, instead of the total variance ξ_0 , on the axes. And the color scale in this case indicates the density of Gaia stars used in each exposure.

We also note that the kernel optimization of Step 5 reduces ξ_0 by a mean factor of $\approx 1.6 \times$ for *riz*-band exposures (with a wide range of variation), and it is possible that this optimization could be improved.

Figure 6 plots $\xi_+(\lvert\mathbf{x}\rvert)$ averaged over all *riz* exposures, showing that (per expectations) the longer-range correlations are completely eliminated by GPR subtraction. The

correlation length, R_0 , is defined as the angle such that $\xi_+(R_0) = 0.5\xi_0$. For the 206 *riz*-band exposures, the average correlation length of the raw residual field was 5.7 . The average correlation length of the residual field after the GPR model was subtracted was 1.3 . As expected, the post-GPR correlation length is similar to the mean distance between Gaia stars used for training.




Figure 6. The points plot the mean azimuthally averaged $\xi_+(|\mathbf{x}|)$ for 206 *riz*-band exposures. The curves apply a Savitsky–Golay smoothing to the (noisy) measured points separation from $5''$ to $1'.5$. From top to bottom, the curves show the correlated astrometric errors before any GPR subtraction (“Raw”), after subtracting a GPR using the best-fit von Karman turbulence parameters (“Fitted model”), and after subtracting a GPR using an “optimized” kernel chosen to minimize $\xi_0 = \xi_+(\mathbf{x} \rightarrow 0)$. The astrometric variance is greatly reduced from the raw level of $\xi_0 = 119 \text{ mas}^2$ to 19.5 and 12.0 mas^2 by the fitted and optimized kernels, respectively. The average correlation length of the astrometric errors (defined as the point where ξ_+ drops to $\xi_0/2$) is reduced from 5.7 to 1.2 by GPR subtraction.

The average von Karman kernel parameters for the *riz* functions are in rough agreement with a physical model, although it is clear from, e.g., Figure 3 that the single-screen turbulence model is an incomplete description of the astrometric correlation functions. The average aperture-diameter parameter, d , was $2'.5$, which corresponds to the physical 4 m diameter D for a turbulence height of 5.5 km , which is of similar magnitude to the typical height of dominant turbulence (Beckers 1993). At this height the average angular outer-scale parameter of $r_0 = 1'.7$ corresponds to 160 m , also physically reasonable—though this parameter has a weak effect on the kernel or GPR accuracy. The typical wind parameter amplitude of $|\mathbf{w}| = 3'.5$ corresponds to winds of just $\approx 0.1 \text{ m s}^{-1}$, which seems lower than one would expect for tropospheric winds.

There are few clear trends in the values of the post-GPR ξ_0 or in the reduction factor gained by the GPR. From the bottom panel of Figure 5 it is clear that the exposures taken in low-density regions of the Gaia catalog have the lowest (worst) reduction factors. This is to be expected, as density Gaia training data can measure and remove the turbulence to higher spatial frequencies. We also find, not surprisingly, that a lower raw ξ_0 correlates with a lower GPR-subtracted ξ_0 . Beyond this, there are no obvious trends with filter band or other variables. The sizes of ξ_0 both before and after GPR subtraction are highly variable over time, as one finds for related atmospheric turbulence phenomena such as the seeing FWHM. These large “weather” variations could be masking subtler trends with wavelength, air mass, etc., that might emerge upon analysis of a larger number of exposures.

6.1. Test on Eris Orbit

As a test of the GPR astrometric correction, we examine the residuals of an orbit fit to the positions measured by DES for the trans-Neptunian object Eris (Brown et al. 2005). Because Eris moves $\approx 1'.4$ in both R.A. and decl. over the 5 yr span of the DES observations, it is being informed by a continuously

changing set of Gaia stars, and thus the orbit residuals are a good sampling of the accuracy of the GPR. A good orbit also requires a correct absolute astrometric calibration, at least across these few degrees.

Figure 7 plots the residuals to the 27 *griz* observations of Eris in the DES Wide survey, compared to the best-fit orbit as obtained using the algorithms of Bernstein & Khushalani (2000). The error bars in the plot are the expected uncertainties from the quadrature sum of the `ERRAWIN_WORLD` measurement error and the ξ_0 -derived rms turbulence error. The measured pixel positions for the Eris detections are converted to R.A. and decl. using the full static astrometric model as described above, including chromatic terms.

Restricting our consideration to the 19 *riz* exposures, the rms residual to the best-fit orbit is 10.1 mas (per axis) using the raw positions (no turbulence correction) but drops to 5.0 mas after the GPR estimates are subtracted. The χ^2 values and degrees of freedom are 27.9 and $44.7/32$ before and after GPR subtraction, respectively. This demonstrates a clear improvement in astrometric quality. The factor of 2 improvement in rms is less than the $\sqrt{12}$ we might expect from the typical reduction in ξ_0 , because Eris, at $r = 18.5 \text{ mag}$, is not bright enough to be in the fully turbulence-dominated regime for DES once the GPR is applied: the shot-noise errors in its positions vary from 1.7 to 5.2 mas (with an rms value of 2.9 mas). Indeed, once the GPR correction is applied, a star must be near the exposure saturation limit in order to have shot noise well below the turbulence noise.

These measurement accuracies are far better than is typically obtained for minor planets from ground-based observations, especially considering the short (90 s) exposure times. To give a sense of scale, a typical trans-Neptunian object at 40 au distance moves by 4 mas s^{-1} , so shutter-timing corrections are now larger than the $\approx 2 \text{ mas}$ astrometric turbulence error.

7. Summary and Prospects

We have demonstrated that GPs are highly effective for interpolating the stochastic astrometric distortions from a set of known spatial points (Gaia stars) to arbitrary locations in the focal plane. For DES images in *riz* bands, we achieve an average reduction of a factor ≈ 12 in the total astrometric variance ξ_0 ascribable to turbulence. As expected, the GPR is very successful at modeling (and removing) distortion modes at wavelengths longer than the typical $\approx 1'$ spacing between Gaia DR2 stars, and the GPR-subtracted astrometric errors have a correlation length of $\approx 1'$. This achievement is assisted by a new variant of GPR that makes use of the known curl-free nature of the 2D distortion field.

For these 90 s images, this reduces the rms turbulence error to $\approx 2 \text{ mas}$ in each axis, at which point it is subdominant to shot noise in the object centroid for any source with $S/N \lesssim 200$, which requires $>40,000$ photoelectrons to be acquired. There is now thus only a $\approx 1 \text{ mag}$ range of stellar brightness in which the astrometry is dominated by turbulence noise but the stars are not saturating the CCD, so in practical terms the atmospheric turbulence noise has been nearly eliminated.

It is likely that some further improvement is possible by improving the optimization of the von Karman parameters, or by choosing a better form for the kernel. It is already clear from Figure 3 that the observed $\xi_+(\mathbf{x})$ function has qualitative aspects not reproduced by the von Karman model, such as different major axes at different scales, which could result from




Figure 7. Residuals of the observed R.A./decl. positions from GPR-corrected DES measurements of the trans-Neptunian object Eris to the best-fitting orbit are plotted. The observations span 5 yr, with the relative times marked below each point pair. The rms error of the *r_{iz}*-band observations is 5 mas, as marked by the shaded region, consistent with the estimated shot noise plus residual turbulence ξ_0 . The rms residual was 10 mas before the GPR correction was made, demonstrating the success of the GPR method.

multiple layers of turbulence. Such behavior has been observed in data from the Canada–France–Hawaii Telescope, and it is worth investigating whether using the measured ξ_+ directly, rather than a model fit, is generically better (P. F. Leget, private communication). We have not developed this yet, because of the possibility that a noisy measured ξ_+ could lead to non-positive-definite covariance matrices.

Lubow et al. (2021) have applied a somewhat simpler approach to removing astrometric errors from the PS1 catalogs by referencing each object to the mean of the nearest 33 Gaia DR2 stars detected in the PS1 images. This would be equivalent to GPR with a flat-top circular kernel of radius equal to the distance of the 33rd-nearest star (reported as $\approx 1'$ at median). The GPR method can be expected to perform better through a weighting of the Gaia reference stars that is optimized for the spectrum of the distortions and the measurement noise of individual references. On the other hand, the GPR is computationally slower since it uses the full set of Gaia references to derive and apply these optimized weights. A direct performance comparison between the two results is not possible given the different telescopes and observing conditions of the surveys, and because Lubow et al. (2021) do not separate turbulence errors from shot noise in their statistics. We do see from their Figure 9, however, that for bright stars ($i \approx 17$ mag), where shot-noise contributions should be low, their correction reduces median 2D astrometric residuals by about a factor 2, whereas we obtain $\sqrt{12} \approx 3.5\times$ reduction in turbulence noise, perhaps as a result of the more optimal interpolation.

Dramatic improvement may be possible by creating a denser training set than Gaia DR2 provides. Future releases of Gaia will help in this regard. But a more powerful means to

bootstrap a denser training sample is outlined in the [Appendix](#): when we require that non-Gaia stars’ motion over time fits the standard parallax/proper-motion model, their 5D solutions become increasingly constrained with more observing epochs, and the exposure-by-exposure residuals to these 5D solutions become useful for constraining the turbulence patterns of individual exposures. The [Appendix](#) gives the mathematical solution for a feasible scheme to solve for the 5D solutions and the turbulence interpolation for all exposures in a region of sky simultaneously. We will test this in future work. The astrometric system, i.e., the absolute coordinates, proper motion, and parallax zero-points, will remain tied to Gaia, and here again ground-based results will benefit from future Gaia releases.

With the joint turbulence/proper-motion solution, essentially every star with per-exposure shot-noise level lower than the turbulence noise becomes an additional reference point for the turbulence solution. While DES has few stars fainter than the Gaia limit that satisfy this condition, LSST will have many. We thus expect this to be a valuable technique for the survey.

Note that both shot-noise errors and turbulence residuals are expected to decrease as the inverse square root of exposure time T . Thus, for a given observatory and filter band, the magnitude dividing turbulence-dominated from shot-noise-dominated astrometry is independent of T . The GPR technique shifts that magnitude brightward by about 2.5 mag. For longer DECam exposures, both limits drop, and astrometry could become limited by sub-milliarsecond errors in the static astrometric model, including errors in chromatic corrections arising from inaccurate knowledge of the source spectra.

This GPR technique should be directly applicable to the LSST data. The Rubin Observatory’s 8.5 m diameter primary

mirror is larger than the Blanco’s 4 m primary, which will lower the expected level of ξ_+ in similar atmospheric conditions, but the shorter total exposure time (30 s) of nominal LSST visits will increase the per-exposure ξ_+ , so we might expect values similar to those we find for DES. Rubin’s greater aperture means, however, that shot noise is lower than for DES, which means that more stars—including many beyond Gaia’s magnitude limit—will be in the turbulence-dominated regime and thus benefit from the GPR subtraction. This should allow LSST data to substantially surpass the initial requirements and goals for its astrometric performance and subsequent science.

We thank Pierre-François L  get for very useful conversations about anisotropic GPs and atmospheric turbulence at other telescopes. University of Pennsylvania authors have been supported in this work by grants AST-1515804, AST-1615555, and AST-2009210 from the National Science Foundation and grant DE-SC0007901 from the Department of Energy.

Funding for the DES Projects has been provided by the U.S. Department of Energy, the U.S. National Science Foundation, the Ministry of Science and Education of Spain, the Science and Technology Facilities Council of the United Kingdom, the Higher Education Funding Council for England, the National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign, the Kavli Institute of Cosmological Physics at the University of Chicago, the Center for Cosmology and Astro-Particle Physics at The Ohio State University, the Mitchell Institute for Fundamental Physics and Astronomy at Texas A&M University, Financiadora de Estudos e Projetos, Funda  o Carlos Chagas Filho de Amparo    Pesquisa do Estado do Rio de Janeiro, Conselho Nacional de Desenvolvimento Cient  fico e Tecnol  gico and the Minist  rio da Ci  ncia, Tecnologia e Inova  o, the Deutsche Forschungsgemeinschaft and the Collaborating Institutions in the Dark Energy Survey.

The Collaborating Institutions are Argonne National Laboratory, the University of California at Santa Cruz, the University of Cambridge, Centro de Investigaciones Energ  ticas, Medioambientales y Tecnol  gicas-Madrid, the University of Chicago, University College London, the DES-Brazil Consortium, the University of Edinburgh, the Eidgen  ssische Technische Hochschule (ETH) Z  rich, Fermi National Accelerator Laboratory, the University of Illinois at Urbana-Champaign, the Institut de Ci  ncies de l’Espai (IEEC/CSIC), the Institut de F  sica d’Altes Energies, Lawrence Berkeley National Laboratory, the Ludwig-Maximilians Universit  t M  nchen and the associated Excellence Cluster Universe, the University of Michigan, NFS’s NOIRLab, the University of Nottingham, The Ohio State University, the University of Pennsylvania, the University of Portsmouth, SLAC National Accelerator Laboratory, Stanford University, the University of Sussex, Texas A&M University, and the OzDES Membership Consortium.

Based in part on observations at Cerro Tololo Inter-American Observatory at NSF’s NOIRLab (NOIRLab Prop. ID 2012B-0001; PI: J. Frieman), which is managed by the Association of Universities for Research in Astronomy (AURA) under a cooperative agreement with the National Science Foundation.

The DES data management system is supported by the National Science Foundation under grant Nos. AST-1138766 and AST-1536171. The DES participants from Spanish institutions are partially supported by MICINN under grants ESP2017-89838,

PGC2018-094773, PGC2018-102021, SEV-2016-0588, SEV-2016-0597, and MDM-2015-0509, some of which include ERDF funds from the European Union. IFAE is partially funded by the CERCA program of the Generalitat de Catalunya. Research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Program (FP7/2007-2013) including ERC grant agreements 240672, 291329, and 306478. We acknowledge support from the Brazilian Instituto Nacional de Ci  ncia e Tecnologia (INCT) do e-Universo (CNPq grant 465376/2014-2).

This manuscript has been authored by Fermi Research Alliance LLC under contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

Software: The software developed in this work will be made public shortly after the publication. This work made use of the following public codes: NUMPY (Van Der Walt et al. 2011), SCIPY (Virtanen et al. 2020), ASTROPY (The Astropy Collaboration et al. 2018), MATPLOTLIB (Hunter 2007), IPYTHON (P  rez & Granger 2007), TREECORR (Jarvis 2015), EASYACCESS (Carrasco Kind et al. 2019), WCSFIT, and PIXMAPPY (Bernstein et al. 2017).

Appendix Simultaneous Inference of Turbulence and 5D Stellar Parameters

The positions of isolated stars on the sky are expected to follow the standard five-parameter model:

$$\mathbf{x}_{i\mu}^{\text{true}} = \mathbf{x}_i^0 + t_\mu \dot{\mathbf{x}}_i - \mathbf{x}_{\perp\mu}^E \varpi_i. \quad (\text{A1})$$

$$= \mathbf{M}_\mu \mathbf{s}_i. \quad (\text{A2})$$

Here t_μ and $\mathbf{x}_{\perp\mu}^E$ are the date of exposure μ and the projection of the barycentric observatory position onto the line of sight for the exposure. The five parameters for star i are the position, proper motion, and parallax $\mathbf{s}_i = \{x_i^0, y_i^0, \dot{x}_i, \dot{y}_i, \varpi_i\}$. The goal of this section is to develop a method for extracting $\mathbf{s} = \{\mathbf{s}_1, \mathbf{s}_2, \dots\}$ from the observed positions $\mathbf{x}_{i\mu}^{\text{obs}}$ of each star in each exposure.⁵⁵ The data model is

$$\mathbf{x}_{i\mu}^{\text{obs}} = \mathbf{M}_\mu \mathbf{s}_i + \mathbf{u}_{i\mu}. \quad (\text{A3})$$

We will always assume that \mathbf{x}^{obs} has been mapped from pixel coordinates to sky coordinates using the best available static instrument model, and \mathbf{u} is a stochastic, zero-mean error term that includes curl-free atmospheric turbulence, shot noise, and any errors in the instrumental model.

The method of Section 3.2 gives a straightforward procedure by which a GPR trained on the Gaia star images in exposure μ yields an estimator $\hat{\mathbf{u}}_{i\mu}$ for each individual image of a non-Gaia star. One could then fit the set of observed $\mathbf{x}_{i\mu}^{\text{obs}} - \hat{\mathbf{u}}_{i\mu}$ to the model (A3) to estimate \mathbf{s}_i . It is the case, however, that as more exposures are taken and the \mathbf{s}_i becomes better known, this star can begin to inform the training of the GPR for each exposure, improving the estimator $\hat{\mathbf{u}}$ for other stars in its vicinity, and in turn improving their 5D solutions. We therefore explore the possibility of an estimator for \mathbf{s} that considers the data from all exposures simultaneously.

To do so, we continue the assumption that the displacement field \mathbf{u}_μ is the sum of a turbulence contribution, which is the

⁵⁵ A given star need not be observed in all exposures.

gradient of a Gaussian random field with known power spectrum, and a shot-noise contribution, which does not correlate between stellar observations. We further assume that there is no correlation of either component between distinct exposures—and each exposure can have its own correlation function K_μ , determined by hyperparameter optimization on the Gaia stars, as described above.

In this case the probability of the observations is a multivariate Gaussian that separates between exposures:

$$p(\{\mathbf{x}_\mu^{\text{obs}}\} | s) = \prod_\mu |2\pi\mathbf{K}_\mu|^{-1/2} \times \exp\left[-\frac{1}{2}(\mathbf{x}_\mu^{\text{obs}} - \mathbf{M}_\mu s)^T \mathbf{K}_\mu^{-1} (\mathbf{x}_\mu^{\text{obs}} - \mathbf{M}_\mu s)\right], \quad (\text{A4})$$

where we have used the data model in Equation (A3) and defined

$$\mathbf{K}_{\mu,ij} = K_\mu(\mathbf{x}_{i\mu} - \mathbf{x}_{j\mu}) + n_{i\mu}\delta_{ij}. \quad (\text{A5})$$

Also, $\mathbf{x}_\mu^{\text{obs}}$ is the concatenation of all the observed positions on exposure μ .

Bayes's theorem, assuming independent Gaussian priors on each star's parameters, gives

$$p(s|\{\mathbf{x}_\mu^{\text{obs}}\}) \propto p(\{\mathbf{x}_\mu^{\text{obs}}\} | s) \times \prod_i \mathcal{N}(s_i^p, \mathbf{K}_i^p). \quad (\text{A6})$$

The prior can include Gaia measurements of the star's s_i , if available. It is also advisable to place a weak prior on the parallax for non-Gaia stars to avoid numerical instabilities for stars that have been observed at the same limited range of dates each year. The observing cadence of DES is not designed for parallax measurements and can produce such degeneracies.

By substituting Equation (A4) into this equation and concatenating the stellar parameters and their priors into $s^p = \{s_1, s_2, \dots\}$ and $\mathbf{K}^p = \text{diag}(\{\mathbf{K}_i^p\})$, we obtain the log posterior for the stellar parameters:

$$-2 \log p(s|\{\mathbf{x}_\mu^{\text{obs}}\}, s^p, \mathbf{K}_p) = (\text{const}) + (s - s^p)^T \mathbf{K}_p^{-1} (s - s^p) \quad (\text{A7})$$

$$+ (\mathbf{x}_\mu^{\text{obs}} - \mathbf{M}_\mu s)^T \mathbf{K}_\mu^{-1} (\mathbf{x}_\mu^{\text{obs}} - \mathbf{M}_\mu s) \quad (\text{A8})$$

$$\Rightarrow s|\{\mathbf{x}_\mu^{\text{obs}}\}, s^p, \mathbf{K}_p \sim \mathcal{N}(\hat{s}, \mathbf{C}_s), \quad (\text{A9})$$

$$\mathbf{C}_s = \left[\mathbf{K}_p^{-1} + \sum_\mu \mathbf{M}_\mu^T \mathbf{K}_\mu^{-1} \mathbf{M}_\mu \right]^{-1}, \quad (\text{A10})$$

$$\hat{s} = \mathbf{C}_s \left[\mathbf{K}_p^{-1} s^p + \sum_\mu \mathbf{M}_\mu^T \mathbf{K}_\mu^{-1} \mathbf{x}_\mu^{\text{obs}} \right]. \quad (\text{A11})$$

The last line gives the maximum-posterior estimate of all the stars' 5D properties. This joint solution is computationally feasible; GPRs are rate limited by Cholesky matrix factorizations. If there are N_s stars appearing in N_e exposures, then each of the \mathbf{K}_μ are $2N_s \times 2N_s$ square, so the sum inside the brackets of Equation (A10) requires $O(8N_e N_s^3)$ operations. The \hat{s} vector has $5N_s$ elements, so the factorization needed to obtain it in Equation (A11) requires $O(125N_s^3)$ operations, which in most cases ($N_e > 15$) is lower than the cost of the per-exposure inversions.

























It is clearly infeasible to execute the solution in one shot over the full DES survey, with $N_s \sim 10^8$ and $N_e \sim 10^5$. There is little

lost, however, in dividing the survey into regions and solving each independently. Each region's solution is informed by stars within a few atmospheric correlation lengths of its targeted boundaries, which should be included in the solution. But more distant stellar measures do not improve the target region and need not be included. Thus, the computation is best managed by dividing the survey into $\approx 1^\circ$ regions, with $N_s \sim 10^5$, $N_e \sim 10^2$.

The computational cost of the wholistic solution can be compared to the cost of optimizing the hyperparameters of \mathbf{K}_μ for each exposure. The optimizer must reinvert the $(2 \times N_G)^2$ matrix many times to maximize the log marginal likelihood, with N_G being the number of Gaia stars in the full exposure, whereas the wholistic solution requires a single $(2 \times N_s)^2$ inversion, where N_s is the total number of stars of interest in the region. In most cases the wholistic solution will take comparable or less time than completing the per-exposure kernel tuning.

The only distinction between Gaia detections and DES-only stars in Equation (A11) is that the former have stronger priors. The solution uses all of the stars to constrain the turbulence field. In practice, one will also want to interpolate the turbulence field to the locations of sources in individual exposures that do not have 5D solutions (e.g., minor planets and other transients). As with a simple GPR, it is straightforward to perform this interpolation, and the computational cost is much lower than for the Cholesky inversion of \mathbf{C}_s .

ORCID iDs

W. F. Fortino  <https://orcid.org/0000-0001-7559-7890>
 G. M. Bernstein  <https://orcid.org/0000-0002-8613-8259>
 P. H. Bernardinelli  <https://orcid.org/0000-0003-0743-9422>
 J. Annis  <https://orcid.org/0000-0002-0609-3987>
 J. Carretero  <https://orcid.org/0000-0002-3130-0204>
 A. Drlica-Wagner  <https://orcid.org/0000-0001-8251-933X>
 K. Eckert  <https://orcid.org/0000-0002-1407-4700>
 E. Gaztañaga  <https://orcid.org/0000-0001-9632-0815>
 D. W. Gerdes  <https://orcid.org/0000-0001-6942-2736>
 R. A. Gruendl  <https://orcid.org/0000-0002-4588-6517>
 J. Gschwend  <https://orcid.org/0000-0003-3023-8362>
 S. R. Hinton  <https://orcid.org/0000-0003-2071-9349>
 D. L. Hollowood  <https://orcid.org/0000-0002-9369-4157>
 N. Kuropatkin  <https://orcid.org/0000-0003-2511-0946>
 F. Menanteau  <https://orcid.org/0000-0002-1372-2534>
 R. Miquel  <https://orcid.org/0000-0002-6610-4836>
 A. Palmese  <https://orcid.org/0000-0002-6011-0530>
 A. Roodman  <https://orcid.org/0000-0001-5326-3486>
 E. S. Rykoff  <https://orcid.org/0000-0001-9376-3135>
 I. Sevilla-Noarbe  <https://orcid.org/0000-0002-1831-1953>
 G. Tarle  <https://orcid.org/0000-0003-1704-0781>
 D. L. Tucker  <https://orcid.org/0000-0001-7211-5729>
 A. R. Walker  <https://orcid.org/0000-0002-7123-8943>
 J. Weller  <https://orcid.org/0000-0002-8282-2010>

References

- Anderson, J., Bedin, L. R., Piotto, G., Yadav, R. S., & Bellini, A. 2006, *A&A*, **454**, 1029
 The Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123
 Beckers, J. M. 1993, *ARA&A*, **31**, 13
 Bernstein, G., & Khushalani, B. 2000, *AJ*, **120**, 3323
 Bernstein, G. M., Armstrong, R., Plazas, A. A., et al. 2017, *PASP*, **129**, 074503
 Bertin, E., & Arnouts, S. 1996, *A&AS*, **117**, 393

- Bouy, H., Bertin, E., Moraux, E., et al. 2013, *A&A*, **554**, A101
- Brown, M. E., Trujillo, C. A., & Rabinowitz, D. L. 2005, *ApJL*, **635**, L97
- Carrasco Kind, M., Drlica-Wagner, A., Koziol, A. M. G., & Petrivick, D. 2019, *J. Open Source Softw.*, **4**, 1022
- Diehl, H. T., Abbott, T. M. C., Annis, J., et al. 2014, *Proc. SPIE*, **9149**, 91490V
- Flaugher, B., Diehl, H. T., Honscheid, K., et al. 2015, *AJ*, **150**, 150
- Gaia Collaboration, Brown, A. G. A., Vallenari, A., et al. 2018, *A&A*, **616**, A1
- Han, I., & Gatewood, G. D. 1995, *PASP*, **107**, 399
- Hunter, J. D. 2007, *CSE*, **9**, 90
- Jarvis, M. 2015, TreeCorr: Two-point correlation functions, Astrophysics Source Code Library, ascl:1508.007
- Lazorenko, P. F., & Lazorenko, G. A. 2004, *A&A*, **427**, 1127
- Lazorenko, P. F., Mayor, M., Dominik, M., et al. 2009, *A&A*, **505**, 903
- Lindgren, L. 1980, *A&A*, **89**, 41
- Lubow, S. H., White, R. L., & Shiao, B. 2021, *AJ*, **161**, 6
- Morganson, E., Gruendl, R. A., Menanteau, F., et al. 2018, *PASP*, **130**, 074501
- Pérez, F., & Granger, B. E. 2007, *CSE*, **9**, 21
- Perryman, M. A. C., Lindgren, L., Kovalevsky, J., et al. 1997, *A&A*, **500**, 501
- Platais, I., Kozhurina-Platais, V., Girard, T. M., et al. 2002, *AJ*, **124**, 601
- Van Der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, *CSE*, **13**, 22
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261
- Zacharias, N. 1996, *PASP*, **108**, 1135