

Analytical and statistical models for laboratory and astrophysical precision measurements

A dissertation presented

by

Nicholas Ryan Langellier

to

The Department of Physics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Physics

Harvard University

Cambridge, Massachusetts

August 2020



Photo: N. R. L.

© 2020 — Nicholas Ryan Langellier

All rights reserved.

Dissertation Advisor:
Ronald Walsworth

Author:
Nicholas Ryan Langellier

Analytical and statistical models for laboratory and astrophysical precision measurements

Abstract

The first detection of gravitational waves (GW) by the Laser Interferometer Gravitational wave Detector (LIGO) ushered in a new era of cosmic exploration and with it came a plethora of new ideas and theories. One such idea includes a proposal to use space-based optical lattice atomic clocks as GW detectors. We propose a GW detector comprised of two free falling strontium-87 optical lattice clocks locked together by a single optical baseline. The two clocks will measure Doppler shifts induced in the two satellites by an incident GW in the transverse direction. Using a series of dynamical decoupling (DD) control sequences, we outline a frequency-tunable, narrowband sensor capable of detecting coherent GW sources such as binary inspirals and mergers of black holes and neutron stars. Furthermore, the tunability of the DD sequences allow for the continuous observation of such sources, bridging the gap between space-based detectors such as Laser Interferometer Space Antenna (LISA) and ground-based detectors.

Another important area of research in astrophysics is the search for habitable exoplanets. However, radial velocity (RV) searches for Earth-mass exoplanets in the habitable zone around Sun-like stars are limited by the effects of stellar variability on the host star. In particular, suppression of convective blueshift and brightness inhomogeneities due to photospheric faculae/plage and starspots are the dominant contribution to the variability of such stellar RVs. Gaussian process (GP) regression is a powerful tool for modeling these quasi-periodic variations. We investigate the limits of this technique using 800 days of RVs from the solar telescope on the HARPS-N spectrograph. These data

provide a well-sampled time series of stellar RV variations. Into this data set, we inject Keplerian signals with periods between 100 and 500 days and amplitudes between 0.6 and 2.4 m s⁻¹. We use GP regression to fit the resulting RVs and determine the statistical significance of recovered periods and amplitudes. We then generate synthetic RVs with the same covariance properties as the solar data to determine a lower bound on the observational baseline necessary to detect low-mass planets in Venus-like orbits around a Sun-like star. Our simulations show that discovering such planets using current-generation spectrographs and GP regression will require more than 12 years of densely sampled RV observations. Furthermore, even with a *perfect* model of stellar variability, discovering a true exo-Venus with current instruments would take over 15 years. Therefore, next-generation spectrographs and better models of stellar variability are required for detection of such planets.

Because the Sun is not a point source, corrections must be made to the solar telescope RVs that are not necessary for normal exoplanet searches. One such correction is the differential extinction that the top and bottom halves of the Sun experience as they traverse through the Earth's atmosphere. As the top half of the Sun will traverse less of the atmosphere than the bottom half, it will experience less atmospheric scattering and absorption and thus will appear brighter relative to the bottom half. In addition, the rotation axis of the Sun is in general not ideally aligned and consequently the red shifted hemisphere will experience a different amount of extinction than the blue shifted hemisphere, resulting in a systematic RV shift that needs correction. We derive here the correction needed as a function of wavelength and fit this expression to the solar telescope RV data. We show that insufficient SNR exists to make a wavelength dependent correction. Thus we derive the wavelength independent differential atmospheric extinction correction used in the GP regression analysis described above.

Finally, we demonstrate the usefulness of statistical modeling in the world of precision measurements using quantum defects in diamond. The field of integrated circuit (IC) security is interested in non-invasive measurements of IC functionality. We employ the use of substitutional nitrogen-vacancy (NV) defects in the diamond crystal lattice as a high sensitivity, high spatial resolution magnetic field imager. This quantum diamond microscope (QDM) measures the magnetic fields emanating from IC circuit activity and produces a high resolution magnetic field image over a wide field of

view (4 by 4 mm). The device under test is chosen as a Xylinx, Artix-7 field programmable gate array (FPGA), which is programmed to activate a user-controlled number of ring oscillators (RO). We take over 1000 magnetic field images of varying numbers of ROs, and use machine learning to automatically predict the number of ROs active given the resulting images. We first use principal component analysis (PCA) as a dimensionality reduction tool, reducing the 10^5 pixels to just 9 PCA scores. These scores are then fed into a support vector machine (SVM) classifier to predict the number of ROs for each image. We report a prediction accuracy of 89% overall with perfect prediction for large numbers of ROs. We repeat this process for an FPGA that has had the plastic casing etched away in order to place the diamond as close to the silicon die as possible. We show that the improved spatial resolution of the magnetic fields leads to perfect prediction accuracy for even single ROs. These results suggest the use of QDMs and machine learning as a promising way forward in the field IC security.

Contents

Abstract	iii
Table of Contents	vi
List of Figures	viii
List of Tables	x
Citations to Previously Published Work	xi
Acknowledgments	xii
Preface	xvi
Co-worker Contributions	xvii
1 Gravitational wave detection with optical lattice atomic clocks	1
1.1 Introduction	1
1.2 Sensing gravitational waves using optical lattice atomic clocks	5
1.3 Derivation of the effective Doppler shift induced by a passing gravitational wave	7
1.4 Outlook	10
2 Detection limits of low-mass, long-period exoplanets using Gaussian processes and the HARPS-N solar telescope	11
2.1 Introduction	11
2.2 Methods	12
2.2.1 Data	12
2.2.2 Gaussian process kernel	13
2.2.3 Fitting the S index	15
2.2.4 Fitting the RVs	18
2.3 Sensitivity map	19
2.3.1 Synthetic planet model	19
2.3.2 Retrieval of injected signals	20
2.4 Synthetic RVs	23
2.4.1 Observing schedule	23
2.4.2 Observation times	23
2.4.3 Synthesized RVs	24
2.5 Discussion and outlook	25
3 Solar atmospheric differential extinction corrections	28
3.1 Introduction	28
3.2 Derivation of differential extinction	30
3.2.1 Limb darkening	30
3.2.2 Radial velocity term	32
3.2.3 Atmospheric extinction	34

3.2.4	Calculation of the angle ϕ	37
3.3	Fitting the RV data	38
3.4	Daily extinction correction used in the RV data	39
4	Magnetic field fingerprinting of integrated circuit activity with a quantum diamond microscope	46
4.1	Introduction	46
4.2	Experimental design	50
4.2.1	QDM experimental setup	50
4.2.2	Integrated circuit preparation, control, and layout	52
4.2.3	Experimental protocol and data analysis	54
4.3	Data for low number active RO states	56
4.4	Machine learning analysis	57
4.4.1	Data preprocessing	57
4.4.2	QDM image dimensionality reduction	59
4.4.3	Integrated circuit activity state classification	63
4.4.4	Classification results	66
4.5	Extended dataset	68
	Appendices	73
	A MCMC samples	74
	B Probing local dark matter density using stellar accelerations	90
B.1	Introduction	90
B.2	Results	92
B.3	Time series analysis using Gaussian processes	94
	Bibliography	97

List of Figures

1.1	Proposed gravitational wave detector	3
1.2	Incident gravitational wave diagram	8
2.1	Gaussian process fit to HARPS-N data	13
2.2	Posterior distribution of Gaussian process fits	14
2.3	Sensitivity map of exoplanet detection with Gaussian processes	19
2.4	Generation of synthetic radial velocities and statistical significance of fits to these data	22
3.1	Expected extinction coefficient in the visible wavelength regime	30
3.2	Limb darkening diagram	31
3.3	Coordinate rotation diagram	32
3.4	Plane parallel approximation for computing airmass	35
3.5	Diagram showing the misalignment of the rotation of the Sun with the great circle connecting the horizon and zenith	36
3.6	Fit of the wavelength dependent extinction coefficient	39
3.7	RV residuals for the wavelength dependent extinction coefficient fit	40
3.8	Individual components of the extinction coefficient fit	40
4.1	Quantum diamond microscope setup and Nitrogen-vacancy physics	47
4.2	Field programmable gate array images	53
4.3	Example magnetic field images	56
4.4	Principal component analysis and support vector machine classification of magnetic field images	58
4.5	Fraction of variance explained by each principal component	59
4.6	Additional principal component score plots	61
4.7	Relative importance of first 9 principal components	62
4.8	Example diagram explaining support vector machines	63
4.9	Cross validation of the support vector machine regularization hyperparameter, C	66
4.10	Subset of the extended quantum diamond microscope magnetic field dataset	69
4.11	Principal component analysis and support vector machine model performance metrics of the extended dataset	70
A.1	MCMC samples for a synthetic 225 day planet with a semi-amplitude of 1 m/s. The data is synthesized with a baseline of 2 years.	75
A.2	MCMC samples for a synthetic 225 day planet with a semi-amplitude of 1 m/s. The data is synthesized with a baseline of 4 years.	76
A.3	MCMC samples for a synthetic 225 day planet with a semi-amplitude of 1 m/s. The data is synthesized with a baseline of 8 years.	77
A.4	MCMC samples for a synthetic 225 day planet with a semi-amplitude of 1 m/s. The data is synthesized with a baseline of 16 years.	78

A.5	MCMC samples for a synthetic 225 day planet with a semi-amplitude of 1 m/s. The data is synthesized with a baseline of 24 years.	79
A.6	MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.5 m/s. The data is synthesized with a baseline of 2 years.	80
A.7	MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.5 m/s. The data is synthesized with a baseline of 4 years.	81
A.8	MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.5 m/s. The data is synthesized with a baseline of 8 years.	82
A.9	MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.5 m/s. The data is synthesized with a baseline of 16 years.	83
A.10	MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.5 m/s. The data is synthesized with a baseline of 24 years.	84
A.11	MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.1 m/s. The data is synthesized with a baseline of 2 years.	85
A.12	MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.1 m/s. The data is synthesized with a baseline of 4 years.	86
A.13	MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.1 m/s. The data is synthesized with a baseline of 8 years.	87
A.14	MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.1 m/s. The data is synthesized with a baseline of 16 years.	88
A.15	MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.1 m/s. The data is synthesized with a baseline of 24 years.	89
B.1	Geometry for observing stellar accelerations in the Milky Way	91
B.2	Gaussian process analysis of the stellar accelerations simulated dataset	93
B.3	Probability density function of fits to a simulated stellar accelerations data with only white noise	94

List of Tables

2.1	Summary of parameters used for Markov chain Monte Carlo sampling for the S index and RV fits of solar telescope and HARPS-N data	16
4.1	Chip state prediction accuracy on the test dataset (ROs = ring oscillators).	66

Citations to Previously Published Work

Parts of this dissertation cover results reported in the following articles:

1. S. Kolkowitz, I. Pikovski, N. Langellier, M. D. Lukin, R. L. Walsworth, J. Ye, “Gravitational wave detection with optical lattice atomic clocks,” *Physical Review D* **94**, 124043 (2016).
2. A. Ravi, N. Langellier, D. F. Phillips, M. Buschmann, B. R. Safdi, and R. L. Walsworth, “Probing Dark Matter Using Precision Measurements of Stellar Accelerations,” *Physical Review Letters* **123**, 091101 (2019).
3. A. Collier Cameron, A. Mortier, D. F. Phillips, X. Dumusque, R. D. Haywood, N. Langellier, C. A. Watson, H. M. Cegla, J. Costes, D. Charbonneau, A. Coffinet, D. W. Latham, M. López-Morales, L. Malavolta, J. Maldonado, G. Micela, T. W. Milbourne, E. Molinari, S. H. Saar, S. Thompson, N. Buchschacher, M. Ceconi, R. Cosentino, A. Ghedina, A. Glenday, M. Gonzalez, C-H. Li, M. Lodi, C. Lovis, F. Pepe, E. Poretti, K. Rice, D. Sasselov, A. Sozzetti, A. Szentgyorgyi, S. Udry, R. L. Walsworth, “Three years of Sun-as-a-star radial-velocity observations on the approach to solar minimum,” *Monthly Notices of the Royal Astronomical Society* **487**, 1 (2019).
4. M. J. Turner, N. Langellier, R. Bainbridge, D. Walters, S. Meesala, T. M. Babinec, P. Kehayias, A. Yacoby, E. Hu, M. Lončar, R. L. Walsworth, E. V. Levine, “Magnetic Field Fingerprinting of Integrated Circuit Activity with a Quantum Diamond Microscope,” *Physical Review Applied* **14**, 014097 (2020).
5. N. Langellier, T. W. Milbourne, D. F. Phillips, R. D. Haywood, S. H. Saar, A. Mortier, L. Malavolta, S. Thompson, A. Collier Cameron, X. Dumusque, H. M. Cegla, D. W. Latham, J. Maldonado, C. A. Watson, N. Buchschacher, M. Ceconi, D. Charbonneau, R. Cosentino, A. Ghedina, M. Gonzalez, C-H. Li, M. Lodi, M. López-Morales, G. Micela, E. Molinari, F. Pepe, E. Poretti, K. Rice, D. Sasselov, A. Sozzetti, S. Udry, R. L. Walsworth, “Detection Limits of Low-mass, Long-period Exoplanets Using Gaussian Processes Applied to HARPS-N Solar RVs,” (in review: *The Astrophysical Journal*; *preprint arXiv:2008.05970* (2020)).

Acknowledgments

Graduate school has been the longest, toughest journey of my life and has seen a plethora of highs and lows. Throughout this adventure, numerous people have been there to pick me up during the lows and celebrate with me during the highs. I would not be crossing the graduate school finish line without these amazing people. As such I would like to take some space to thank everyone who has made this possible.

First, I would like to thank my good friends. To my fellow classmates in undergrad at the University of Illinois at Urbana-Champaign Karen, Felicia, Niki, Mark, Ian and Ian, Matt, Emily, Sean, and many more, thank you for spending countless long nights solving problem sets and occasionally skipping out on homework to go out and have fun. Ben, thank you for being a great roommate and friend and motivating me to always put in 100%. Stephanie, I will always remember answering your father's random physics questions. To my graduate school friends Junjie, Annie, David, Susan, Katya, Andy, Anastasia, Naveen, Charlotte, and many more, thank you for all the countless hours climbing, gaming, enjoying great food, and countless other activities to keep me sane.

Next, to my colleagues and collaborators, thank you for helping me progress professionally as a scientist and academic. Chih-Hao Li and Alex Glenday mentored me as a young grad student. Gabor Furesz taught me invaluable hands-on lab skills such as photonic crystal fiber tapering. Long discussions with Shimon Kolkowitz and Igor Pikovski made the atomic clock gravitational wave proposal a reality. Edlyn Levine and Matthew Turner have made the integrated circuit security project one of the most enjoyable projects I have had the pleasure of working on. All of the Walsworth group members have provided insightful comments over numerous group presentations. The entire HARPS-N collaboration has been an honor to work with. The conference in Göttingen, Germany was by far the best conference I have attended thanks to the wonderful conversations with these collaborators. I would like to specifically point out Raphaele Haywood, Andrew Collier Cameron, Steve Saar, and Annelies Mortier. I could not have done this without your help.

A very special thank you goes out to Lisa Cacciabaudo and Jacob Barandes for being the go-to people with any administrative problems and generally being amazing people who are always there for graduate students when they have just about any problem.

To my office mates Tim Milbourne and Aakash Ravi, thank you for making the workplace so enjoyable that I look forward to going in each day. As I move to the next chapter of my life, I will truly miss working with you and spending far too many hours building and fixing CoffeeBot and solving word puzzles. As they say: “Don’t even do work!”.

I would be remiss to fail to mention my family. To my extended family, thank you for all your support. To my brothers Robert and Jacob, thank you for all the fun we had growing up together and continue to have as adults. Most importantly, thank you Mom and Dad for making this possible. You supported me in any endeavor I pursued, from sports to music to relationships to career choices. You taught me how to succeed and I could not ask for better parents. Thank you.

Another very special thank you goes to David Phillips. You have acted as a second scientific advisor from beginning to end of my grad student life. The ability to create an atmosphere of creativity free of judgment is a skill you have mastered and is a large contributor to my success. Your door is always open and you are always ready to tackle tough discussions, whether they be problems scientifically or personally.

To my thesis committee Cora Dvorkin, Dimitar Sassellov, and Julia Mundy, thank you advising me and always being patient with me. In particular I thoroughly enjoyed working with Dimitar in my early grad student career. You have an ability to bring together a diverse set of expertise in way that always results in a whole that is greater than the sum of its parts.

To my advisor Ron Walsworth, thank you for everything. When I was accepted into the Harvard physics PhD program, I visited your lab and you immediately welcomed me into your scientific family. The number of life lessons you have taught me are endless. You provided me many exciting scientific avenues and allowed me to blaze my own trail at times as well. I have enjoyed working

with you and your group. Thank you.

And finally, the biggest thank you goes out to my fiancée Yuanheng Li. When we first met at the climbing gym, I could not have anticipated how my life would change forever. You taught me how to overcome my fear of writing. You picked me up every time I stressed out about end of PhD problems. And you have always believed in me. I could not ask for a better partner and I cannot wait to spend my life with you. I would not be here today without you. Thank you.

To all the names I have surely missed, thank you. To everyone, thank you again for all your support and encouragement. As this journey comes to an end, I am excited to see where life takes me next, and I hope you will all continue to be there with me, cheering me on as you always have.

To Yuanheng and my parents

Preface

With the development of modern computing power came an era of rapid development in both science and technology. Computers can now perform calculations that would have taken years in mere seconds, allowing innovative creations such as the self-driving car. Consequently, scientific modeling of various phenomena can be accomplished analytically when feasible and statistically when infeasible. This work aims to demonstrate the important union of both analytical and statistical modeling of precision measurements in fields ranging from astrophysics to AMO physics to quantum sensing.

CHAPTER 1 discusses a proposal to use optical lattice atomic clocks as gravitational wave (GW) detectors complementary to future space-based interferometric detectors. Included is an analytical derivation of the fractional frequency shift experienced by the atomic clocks as a GW passes by in the transverse direction of an arm of two clocks on board two free falling satellites. **CHAPTER 2** introduces the radial velocity (RV) method of exoplanet detection and the statistical modeling technique of Gaussian process (GP) regression. Synthetic planetary signals are introduced into real RV data taken of the Sun, and GP regression is used to investigate the power of this technique to discover habitable planets like Earth in other solar systems. These solar RVs require corrections because the top half of the Sun experiences less scattering and absorption in the Earth's atmosphere than the bottom half. **CHAPTER 3** derives this correction analytically and is used to fit the solar RV data. **CHAPTER 4** explores the use of another statistical model in the area of integrated circuit (IC) security. Images of magnetic fields emanating from ICs, acquired with a quantum diamond microscope (QDM), are used to classify various operating states of the IC using machine learning techniques such as principal component analysis and support vector machines.

Together, the examples in this work show the power of various modeling techniques in multiple fields of fundamental and applied science.

Co-worker Contributions

CHAPTER 1

Shimon Kolkowitz (University of Wisconsin) carried out much of the AMO theory. Igor Pikovski (Stevens Institute of Technology) performed the general relativity theory. I contributed to the idea conceptualization, derived the fractional frequency difference of the two satellites, and helped with the alternate technology comparison. Ron Walsworth (now at University of Maryland), Mikhail Lukin (Harvard University), and Jun Ye (University of Colorado Boulder) provided valuable insights and technical advice.

CHAPTER 2

I performed all of the simulations and analysis and wrote much of the manuscript. Timothy Milbourne (Harvard University) helped write some of the manuscript and provided helpful technical advice. David Phillips (CfA) helped edit the entire manuscript and provided scientific advice throughout. Ron Walsworth supervised the project. The remaining authors provided valuable feedback on the manuscript and helped build and maintain the telescope hardware and software.

CHAPTER 3

I performed the wavelength dependent extinction correction derivation and the fitting to this function. Andrew Collier Cameron (University of St. Andrews) performed the original derivation of the wavelength independent extinction correction. Annelies Mortier (Kavli Institute) and I helped with derivations of higher order terms of this derivation.

CHAPTER 4

Ron Walsworth and Edlyn Levine (MITRE) conceived of the idea to use QDMs to measure magnetic fields of ICs. Marko Lončar (Harvard University), Amir Yacoby (Harvard University), and Evelyn Hu (Harvard University) provided valuable technical advice throughout. Rachel Bainbridge (MITRE) and Dan Walters (MITRE) provided useful feedback on FPGA usage. Srujan Meesala (CalTech) and Thomas Babinec (Harvard University) took the detailed images of the FPGAs and helped with the decapping of the FPGA. Matthew Turner (Harvard University) built the QDM, took all the data, and performed the analysis to convert the raw data into magnetic fields. I performed the machine learning analysis that predicts IC activity from magnetic field images. Matthew, Edlyn, Pauli Kehayias (Sandia), and I wrote most of the manuscript, with editorial input from the other authors.

Gravitational wave detection with optical lattice atomic clocks

1.1 Introduction

The first direct detections of gravitational waves (GWs) by the Laser Interferometer Gravitational-Wave Observatory (LIGO) [1, 2] heralds the dawn of a new era of astrophysics. The culmination of a century-long search [1–10], GW detection is now emerging as a new tool with which to study the universe, illuminating previously invisible astrophysical phenomena. In parallel, the developments of laser cooling and the laser frequency comb have given rise to optical atomic clocks with accuracies and stabilities at the 10^{-18} level [11–15]. As clock precision continues to improve, there is growing interest in the prospect of using optical atomic clocks for GW detection [10, 16, 17]. We outline a proposal for a new GW detector based on Doppler shift measurements between two spacecraft containing optical lattice atomic clocks linked over a single optical baseline. This detector offers broad tunability of narrowband sensitivity in the mHz - Hz frequency range. As GW astronomy matures, such a detector can therefore serve as a different type of observatory for gravitational waves that can be complementary to existing concepts, much like there are applications for both large and narrow field-of-view telescopes in electromagnetic astronomy. We analyze the prospects for GW detection and characterization using our clock based scheme. We highlight new and complementary GW

measurement capabilities provided by space-based optical atomic clocks, and discuss the prospects for integrating our scheme with existing proposals.

While there is little doubt that LIGO and other terrestrial detectors will observe numerous additional GW events in the coming years, terrestrial detectors are only sensitive to GWs with frequencies above ~ 10 Hz, due to seismic and Newtonian noise [1, 3, 18, 19]. The desire to observe a wider range of astrophysical phenomena over longer length and time scales has motivated proposals of larger scale, space-based GW detectors [16–23]. There are a wide variety of existing and proposed techniques [1–10, 24], all of which rely on the same GW effect, namely the periodic change in proper distance between two points in space [25]. This effect results in modulation of the arrival times of photons sent over an electromagnetic baseline, which corresponds to effective changes in relative position and velocity. The differences between the various techniques lie in the detection methods, the physical quantity that is being locally measured, and the susceptibility to different noise sources, making particular schemes better suited for specific GW frequency ranges.

The existing and proposed space-based GW detectors can be broadly classified in two categories. The first are optical interferometric detectors analogous to LIGO in space, such as the proposed Laser Interferometer Space Antenna (LISA) [18] and Evolved-LISA (eLISA) [19], which would be composed of three spacecraft forming either a two or three arm Michelson interferometer, with roughly equal length arms to reduce susceptibility to laser frequency noise. These GW detectors rely on large photon fluxes to split the optical interference fringe down to the required sensitivities, and detect signals in a broad frequency band determined by the detector arm length and residual acceleration noise of the satellites [18, 19].

The second class of space-based GW detectors rely on stable internal frequency references, such as Doppler tracking of distant spacecraft [7–10]. These detectors search for changes in the frequency of electromagnetic waves due to effective Doppler shifts arising from passing GWs. Doppler tracking of spacecraft has been successfully employed to set the existing limits on milliHertz gravitational wave events [7–10]. Because the sensitivities of this class of detector are generally limited by the stability of the frequency reference rather than the photon flux [10, 26], there is a clear motivation

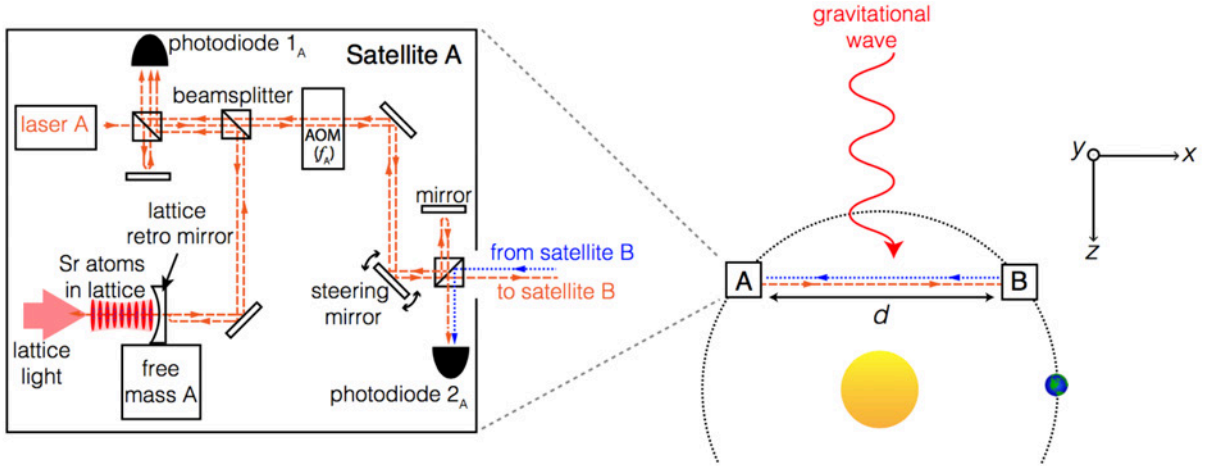


Figure 1.1: Proposed gravitational wave classified in two categories. The first are optical interferometric detector (not to scale). Our detector consists of two identical drag-free satellites, A and B, separated from each other by a distance d along the x -axis. Each satellite contains a free-floating reference mass, an ultra-stable laser, and a strontium optical lattice clock. A mirror is mounted on the free mass and is used to define the standing wave of light forming the optical lattice and confining the Sr atoms. Some of the laser light from satellite A (orange, dashed line) is sent to satellite B. The light first passes through an acousto-optic modulator (AOM) driven at frequency f_A , which offsets the frequency of the light reaching photodiode 2_B in satellite B and enables the phase locking of laser B to laser A through heterodyne detection. Vibrations and thermal drifts of the optics on each satellite can be corrected locally by feeding back on the beat notes at $2f_{A,B}$ on photodiodes $1_{A,B}$. Light from laser B (blue, dotted line) is sent back to satellite A to verify the phase lock, to maintain pointing stability, and to enable operation in the reverse mode, with laser A locked to laser B. A plus-polarized gravitational wave propagating along the z -axis induces relative motion between the two free masses (see section 1.3), which can be detected using a clock comparison measurement protocol. The satellite configuration and orbit shown here is intended only for illustration of the basic concepts of our detector. A more sophisticated orbital pattern could be employed to increase the rate of rotation and sweep the detector pattern over a larger region. Additional satellites and optical links could also be used for improved sensitivity and localization of GW sources.

to improve the internal frequency references used for GW detection through the adoption of atomic physics techniques, such as either atomic interferometry (AI) [20–23], or optical lattice atomic clocks, as described here.

A GW detector composed of two satellites carrying optical lattice atomic clocks and sharing a single laser over an optical link can measure shifts in the rate of optical phase change by comparing the laser frequency to an atomic degree of freedom on both ends of the baseline. Such a detector is therefore similar to the Doppler tracking method of GW detection [7–10, 17], in that it is sensitive to changes in the apparent relative velocities of the reference masses, rather than changes in the apparent relative distance. An important advantage of the atomic clock scheme over other Doppler tracking methods is that one has full control over the frequency references. As a result, synchronized measurement sequences can be applied at both ends of the baseline to cancel laser frequency noise [21, 27–30]: thus the atomic clock technique requires only two spacecraft, not three, and the differential measurement is entirely limited by the internal atomic transition, not the stability of the local oscillator used to probe it [10–15]. Furthermore, dynamical decoupling (DD) control sequences can be applied to the internal states of the atoms [31, 32], extending the range of GW frequencies to which the detector can be maximally sensitive, from milliHertz to tens of Hertz, without requiring any physical changes to the detector. This key feature provides a tunable, narrowband GW detector for tracking evolving GW sources such as inspiraling black hole or neutron star binaries, and bridging the spectral gap between space-borne and terrestrial optical interferometric GW detectors [1–3, 18, 19].

Due to the Doppler-based measurement scheme and high quality factor of atomic clock transitions, the optical power requirements on the link between satellites differ from those of optical interferometer GW detectors, as discussed below. In addition, quantum techniques such as atomic spin-squeezing and entangled states [33–35] offer the potential for future improvements in sensitivity, detection bandwidth, and spectral range. Finally, because optical atomic clocks are currently the most accurate frequency references [11–15], and can provide improved sensitivity to beyond-Standard-Model phenomena that may couple to atomic properties such as mass, charge, and spin [36–39], there is already considerable motivation to develop space-hardy optical clocks, and to integrate them with other proposed GW detectors.

1.2 Sensing gravitational waves using optical lattice atomic clocks

Our proposed GW detector, illustrated in Fig. 1.1, consists of two drag-free satellites in heliocentric orbit (A and B), separated by a length d and connected over a single optical link using conventional optical telescopes. Each satellite contains its own optical lattice atomic clock [11–13, 40], and its own ultra-stable laser [41]. The laser in satellite B is kept phase locked to the light sent from satellite A over the optical link, such that the two lasers function as a single ultra-stable clock laser shared between the two satellites. In each satellite the lattice confining the clock atoms is created using the standing wave formed by retro-reflecting a magic wavelength laser [11–13] off of a mirror mounted on a free-floating reference mass, such that the atoms are strongly confined in the reference frame of the free mass and are therefore in free-fall, despite their confinement. Drag-free masses have been studied in great detail by the LISA collaboration, and this technology is currently undergoing testing and verification in the LISA Pathfinder space mission [42, 43]. The phase of the clock lasers in each satellite is kept referenced to the same mirror using interferometry [44] to cancel out any relative motion of the lasers or optics with respect to the atoms. To cancel the radiation pressure exerted on the free mass by the lattice and clock beams, a set of equal power, counter-propagating lasers are incident on the opposite sides of the free masses. For a 1 kg mass and a 1 W lattice beam, the remaining acceleration noise from the quantum radiation pressure shot noise of the lattice, clock, and compensation beams is far below the GW detector noise floor at frequencies of interest [45].

Operation of the GW detector consists of a synchronous comparison between the two optical lattice atomic clocks. The frequency of laser A is compared to the clock transition in the atoms in satellite A using spectroscopic read-out, such as Ramsey spectroscopy. Synchronization signals are transmitted to satellite B, so that an identical measurement is performed on the atoms in satellite B using laser B, which is phase locked to laser A. Both Ramsey measurements are performed with the same interrogation time T . The Ramsey phases accumulated by the atoms in each satellite are recorded, and can then be compared over a standard communication channel. Because the satellites effectively share a single laser and the two measurements are offset by the time required for the laser light to travel from A to B, any laser frequency noise will be common mode for the two measurements,

resulting in the same additional acquired phase in each clock, and will thus be rejected. This method of laser frequency noise rejection has been previously utilized in optical atomic clocks to cancel laser noise arising from the Dick effect, and thereby achieve the quantum projection noise limit [27, 28, 30]; it has also recently been proposed for use in AI-based GW detectors [21].

A passing plus-polarized GW of strain amplitude h and frequency f_{GW} , propagating along the z -axis perpendicular to the optical link between the satellites, will periodically change the apparent distance between the free masses A and B, as measured by the null geodesic of the optical link. If the light sent from satellite A to B is used as reference clock light, its frequency will experience a Doppler shift that indicates the GW induced effective relative motion of the two satellites. Hence the atoms in satellite B will experience a local oscillator of a different optical frequency than the atoms in satellite A, and will accumulate a different Ramsey phase. When the two clocks are compared, they will appear to have “ticked” at different rates, with the maximum fractional frequency difference between the two clocks given by

$$s \equiv \frac{\delta\nu}{\nu} = h \left| \sin \left(\pi f_{\text{GW}} \frac{d}{c} \right) \right|, \quad (1.1)$$

where c is the speed of light (see section 1.3 for derivation). Note that $s = h$ for the optimal clock spacing $d = \lambda_{\text{GW}}/2$, where $\lambda_{\text{GW}} = c/f_{\text{GW}}$ is the GW wavelength. When compared to an optical interferometric GW detector such as LISA [18, 19], with total optimized arm length d , the fractional frequency difference s between two optimally spaced clocks will be equivalent to the fractional change in differential arm length experienced by the optical interferometer. At GW frequencies other than the optimal frequency the magnitude of the detectable signal is determined by the inherent sensitivity of the specific setup, as captured by the detector’s transfer function $\mathcal{T}(f)$ [10, 46] and susceptibility to noise. The noise floor of optical interferometric detectors is fundamentally limited by white phase noise arising from photon shot noise [19], while the noise floor of the clock detector is dominated by white frequency noise arising from atom projection noise [47]. This fundamental physical difference motivates the present consideration of the former as a detector of changes in phase, and the latter as a detector of changes in frequency, so that the detector transfer functions can be directly compared. We emphasize that while both types

of detector can in principle express their measurement in terms of either phase or frequency, the respective fundamental physical noise floors and signal to noise ratios will be unchanged.

The transfer function $\mathcal{T}_\phi(f)$ for optical interferometric GW detectors such as LISA is frequency independent for GW frequencies below $c/2d$, but scales as $\mathcal{T}_\phi(f) \propto 1/f_{\text{GW}}^2$ at higher frequencies where the photon transit time is longer than a half period of the GW [18, 19]. Because the 1 mHz - 1 Hz frequency range is of primary interest for space-based detectors [18–21], $\mathcal{T}_\phi(f)$ sets a maximum arm length for an optical interferometer on the order of $\sim 1 \times 10^9$ meters. In contrast, because the clock GW detector compares the local laser frequency at the two satellites and is thus only sensitive to the effective relative velocity of the satellites, the transfer function $\mathcal{T}_\nu(f)$ of the clock GW detector scales as $\mathcal{T}_\nu(f) \propto f_{\text{GW}}^2 \times \mathcal{T}_\phi(f)$ due to the time derivative relating position (phase) to velocity (frequency). Therefore, $\mathcal{T}_\nu(f) \propto f_{\text{GW}}^2$ for $f_{\text{GW}} < c/2d$, but is frequency independent at higher frequencies[10]. We are interested in GW frequencies of \sim mHz and above, thus we propose a clock GW detector with a baseline length $d = 5 \times 10^{10}$ m, setting the minimum frequency that can be detected at the detector’s peak sensitivity to be $c/2d \approx 3$ mHz. Note that a LISA-like baseline length of 5×10^9 meters could be used for the clock GW detector without sacrificing sensitivity at GW frequencies above ~ 30 mHz.

1.3 Derivation of the effective Doppler shift induced by a passing gravitational wave

A passing gravitational wave (GW) induces periodic changes in the light travel time between emitter and detector. In this section we derive the magnitude of this effect as a function of the GW amplitude, the orientation between the satellites and the direction of propagation of the GW, and the distance between the clocks. Similar analyses have been performed for proposed detectors that utilize Doppler tracking [48] and pulsar timing [6]. Our detection scheme involves only a one-way link as in the case of pulsar timing, but with full experimental control on both sites for the emission and detection of the signal.




Figure 1.2: A GW incident along the z -axis periodically changes the light travel distance between A and B .

Weak gravitational fields are captured by a perturbed metric $g_{\mu\nu} = \eta_{\mu\nu} + h_{\mu\nu}$, where $\eta_{\mu\nu}$ is the Minkowski metric and $|h_{\mu\nu}| \ll 1$ is a small perturbation. GWs are described in the transverse traceless gauge by the metric:

$$g_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 + h_+ & h_\times & 0 \\ 0 & h_\times & 1 - h_+ & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (1.2)$$

where $h_+(t - z/c)$ and $h_\times(t - z/c)$ correspond to the two polarizations of the wave, which travels in the z -direction. For simplicity we first calculate the effect for a plus-polarized plane wave with $h = h_+ = |h|e^{-i2\pi f(t-z/c)}$, where f is the frequency of the wave and $|h|$ its amplitude (arbitrary polarizations are restored with the substitution $|h| \rightarrow |h_+| \cos(2\psi) + |h_\times| \sin(2\psi)$, where ψ is the polarization angle). The line element for this metric is then

$$ds^2 = -c^2 dt^2 + (1 + h)dx^2 + (1 - h)dy^2 + dz^2. \quad (1.3)$$

We now consider the situation depicted in Fig. 1.2, where a light signal is sent at time t from system

A to system B , which is at a distance d in the $x - z$ -plane. A light-like curve is defined by $ds^2 = 0$. Parameterizing the curve by r with $x = r \sin \theta$, $y = 0$ and $z = r \cos \theta$, the coordinates for the curve become (to lowest order in h):

$$cdt = \left(1 + \frac{1}{2}h \sin^2 \theta\right) dr. \quad (1.4)$$

As the signal is emitted at coordinate time t and travels from A to B in a time $t_1 = t + d/c$ to lowest order in h , it travels an apparent distance

$$D_{AB} = c \int_t^{t_1} dt' = \int_0^d \left(1 + \frac{1}{2}h(1 - \cos^2 \theta)\right) dr, \quad (1.5)$$

where the GW is parameterized by $h = h(t + r/c - r \cos \theta/c)$. In terms of the indefinite integral of the wave, $H(t)$, the above expression becomes

$$D_{AB} = c(t_1 - t) = d + \frac{c}{2}(1 + \cos \theta) \left[H(t) - H\left(t + \frac{d}{c}(1 - \cos \theta)\right) \right] \quad (1.6)$$

In flat space the distance traveled by the light would just be given by d , but the presence of the GW periodically changes the apparent length of the light path. In Doppler tracking techniques, the signal is reflected back to A and measured there. Here, instead, we consider measurement directly on B . The rate of change gives a Doppler shift of the signal $\sigma \equiv \dot{D}_{AB}/c = \Delta\nu/\nu$, where ν is the optical frequency:

$$s = \frac{\Delta\nu}{\nu} = \frac{1 + \cos \theta}{2} \left[h(t) - h\left(t + \frac{d}{c}(1 - \cos \theta)\right) \right]. \quad (1.7)$$

This apparent Doppler shift is the signal to be detected. The effect is maximized for $\theta = \pi/2$, i.e. for the detector aligned perpendicularly to the GW, while the signal disappears for $\theta = 0$, i.e. in the direction of propagation of the GW. Similarly to interferometric detection schemes, the frequency shift is due to transversal motion of test bodies as the GW is passing.

From equation (1.7), we can see that when using a single shared local oscillator to compare two

clocks positioned a distance d apart in the plane ($\theta = \pi/2$) of a passing GW of amplitude $|h|$ and wavelength $\lambda_{\text{GW}} = c/f$, the clocks will appear to “tick” at different rates, with the maximum fractional frequency difference between the two clocks given by

$$s_{max} = |h| \left| \sin \left(\pi \frac{d}{\lambda_{\text{GW}}} \right) \right|. \quad (1.8)$$

Note that the detector is insensitive to GWs with wavelengths that match a multiple of the baseline d .

1.4 Outlook

In order to attain the highest possible sensitivity we consider a next generation strontium-87 optical lattice clock, as this clock has the narrowest demonstrated clock transition linewidth [11, 47]. With expected improvements to these clocks we may achieve a minimum detectable fractional frequency difference of $s_{min} = 1.1 \times 10^{-20}/\sqrt{\text{Hz}}$ which is sufficient for detection of narrowband coherent sources such as compact binary inspirals. In order to maintain sensitivity at high frequencies, we can use a dynamical decoupling (DD) sequence consisting of a Ramsey sequence combine with a series of periodic π -pulses matched to the frequency of the GW [31, 32]. This effectively yields a narrowband detector of GW that is tunable from the mHz to 10s of Hz regime. We can use this tunable detection band to follow inspiral GW as they chirp from the frequency band in which Lisa is sensitive all the way through to where ground-based detectors, such as LIGO, are sensitive. Once an on-going GW event has been detected, the spacing of the π -pulses in the DD detection sequence can be chirped along with the signal to remain optimally sensitive to the particular event throughout its evolution. Thus, atomic clock GW detectors can serve as a complementary detector to space-based optical interferometry methods, possibly operating in parallel to a space-based interferometer such as LISA.

Detection limits of low-mass, long-period exoplanets using Gaussian processes and the HARPS-N solar telescope

2.1 Introduction

State-of-the-art radial velocity (RV) searches for low-mass, long-period exoplanets are limited by signals produced by stellar magnetic variability. An Earth-like or Venus-like planet in orbit around a Sun-like star in its habitable zone induces a reflex RV signal on the order of 10 cm s^{-1} . However, the presence of acoustic oscillations, magnetoconvection, large-scale magnetic structures, and other stellar processes induce RV perturbations that can exceed 1 m s^{-1} (see [49, 50] and references therein). This stellar variability can conceal and even mimic planetary signals in RV surveys, and has resulted in many false detections (e.g., the disproval of CoRoT-7d, [51]; GJ 581d and g, [52]; Alpha Centauri Bb, [53]). Furthermore, these stellar processes act on timescales between minutes and months [54, 55]. For Sun-like stars, the dominant contributions to these intrinsically driven RV perturbations are from the suppression of convective blueshift and brightness inhomogeneities modulated at the rotation period [56, 57]. The wide range of timescales and non-trivial correlations between these processes require a sophisticated statistical framework to decouple stellar activity processes from planetary signals. [58] and [59] study these effects by recovering injected planets of

known properties into real data.

We use state-of-the-art Gaussian process (GP) regression to account for the temporal correlations of rotationally-modulated stellar activity [60, 61]. This GP regression is trained on solar data, as measured by a purpose-built solar telescope feeding the HARPS-N spectrograph [62] operating at the Telescopio Nazionale Galileo (TNG) in the Canary Islands [63, 64]. In section 2.2, we first present the solar data and the GP regression along with the resulting fit on the daily-averaged solar RVs. In section 2.3, we introduce synthetic planets of varying semi-amplitude and orbital period in order to determine the sensitivity of the GP regression to temperate, low-mass planet searches. We conclude in section 2.4 with an analysis of the baseline of RV observations and model assumptions that are necessary for a true exo-Earth detection.

2.2 Methods

2.2.1 Data

We take 5 minute disk averaged exposures of the Sun using the solar telescope and the HARPS-N spectrograph at the TNG, and use a baseline of around 800 days of near-continuous, daytime, solar spectra [63, 64]. For each exposure, the HARPS-N Data Reduction Software (DRS) [65, 66] computes the barycenter-corrected RV with 40 cm s^{-1} single exposure precision and the Mt. Wilson S index [67], a measure of stellar magnetic activity derived from chromospheric re-emission in the core of the singly-ionized Ca H and K line cores [68]. The resulting RVs are further reduced, as described in [69], to remove the RV signatures of the solar system planets, effects of differential extinction across the solar disk in the Earth’s atmosphere (see chapter 3), and other systematic effects due to the Earth’s orbit around the Sun. To mimic the sampling of a typical stellar observing schedule while preserving the exquisite signal-to-noise ratio (SNR) of the solar telescope [64], we compute daily averaged values of each quantity including the RVs, the S index, and the corresponding mean Julian date. While realistic stellar observing schedules would not allow for this level of

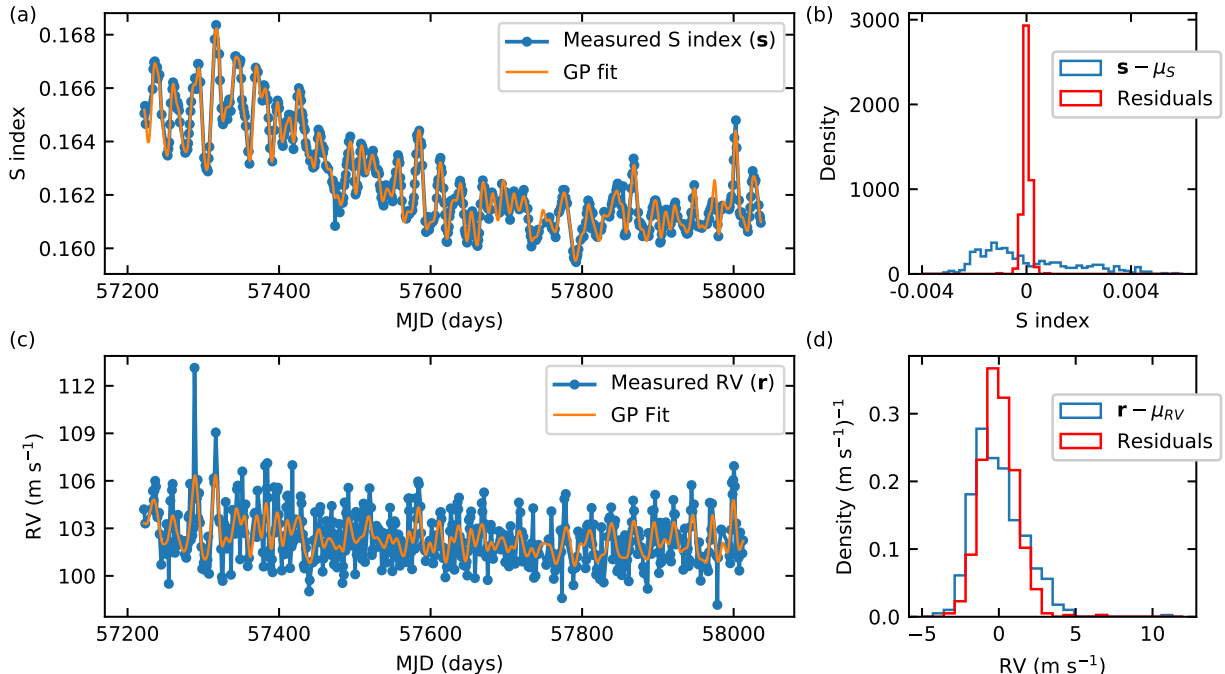


Figure 2.1: GP regression fits to ~ 800 days of solar S index and RV data from the solar telescope and HARPS-N. (a) S index (blue) and GP regression fit (orange). (b) Histogram of mean-subtracted S index values (blue) and histogram of the residuals of the GP regression fit (red). (c) Daily average RVs (blue) and GP regression fit (orange). The mean value represents the HARPS-N instrumental offset and carries no physical meaning. (d) Histogram of mean-subtracted RVs (blue) and histogram of the residuals of the GP regression fit (red). The RMS variation is reduced from 1.65 m s^{-1} to 1.14 m s^{-1} . The relative size of the white noise compared to the correlated noise is higher in the RVs than in the S index, emphasizing the need for a more sophisticated model to account for RV variation.

averaging, which integrates over variability on the minutes and hours timescale, we wish to assess the best-case scenario using an ultra-high SNR dataset.

2.2.2 Gaussian process kernel

Magnetic variability can be the dominant source of variance for stellar RVs of nearby, bright Sun-like and low-mass stars, reaching levels surpassing 1 m s^{-1} [70, 71]. Intrinsic stellar variability introduces correlations into the RV time series that are difficult to model deterministically. GP regressions have emerged as a powerful statistical technique that relaxes the assumption of uncorrelated noise by adding nonzero terms to the off diagonal of the data covariance matrix. Usually a kernel function is

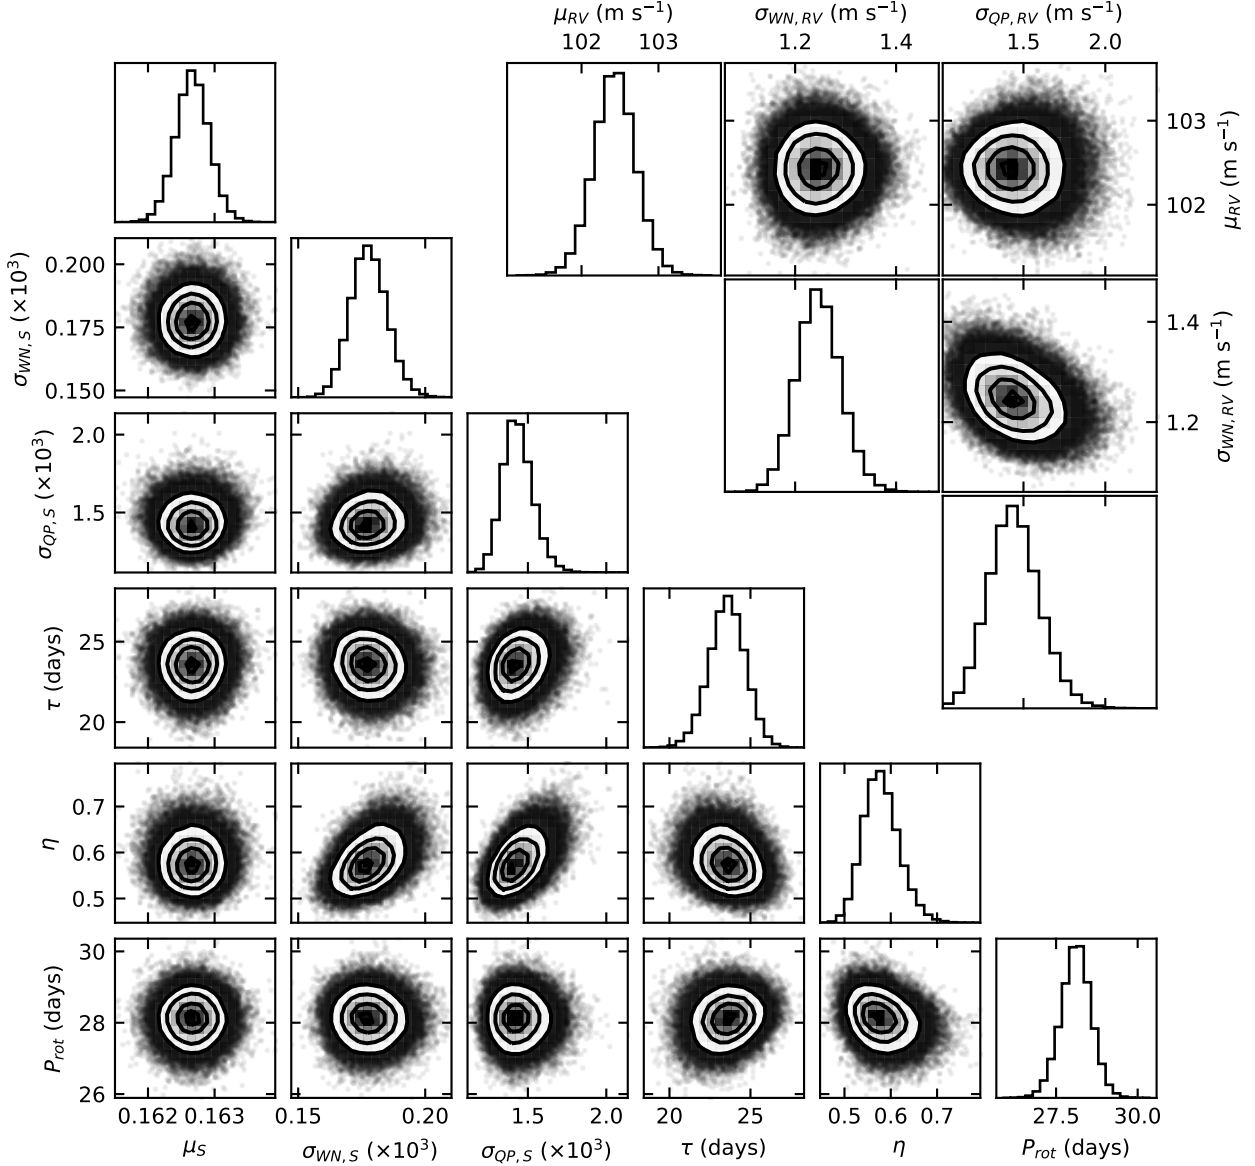


Figure 2.2: Marginalized distributions of the posterior distribution given by the MCMC samples. The lower left triangle shows samples (and contour lines) from the S index fit and the upper right triangle shows samples (and contour lines) from the RV fit. Each parameter is unimodal and only slight correlations exist between parameters. Thus both posterior distributions are well behaved.

chosen to describe these covariances as a function of measurement separation time [51, 61, 72–74].

Since magnetic activity is well described by the S index and is modulated at the rotation period of the star, we fit the S index to a GP with a quasi periodic (QP) covariance kernel function, k_{QP} . This kernel function is chosen heuristically to model the known properties of the magnetic activity, yielding the correlation between two measurements at times t_i and t_j given by

$$k_{\text{QP}}(t_i, t_j) = \exp\left(-\frac{(t_i - t_j)^2}{2\tau^2} - \frac{1}{2\eta^2} \sin^2\left(\frac{\pi(t_i - t_j)}{P_{\text{rot}}}\right)\right) \quad (2.1)$$

where τ is related to the average lifetime of active regions on the Sun, η is related to the average distribution of activity in the photosphere of the Sun, and P_{rot} is approximately the synodic rotation period of the Sun. [75] show this kernel function performs the best among three common GP kernel functions. The GP regression also includes a white noise (WN) term to account for additional uncorrelated noise sources such as instrumental effects. Amplitudes of both the correlated and uncorrelated terms are allowed to vary, leading to a covariance matrix, \mathbf{K}_S , with components

$$K_{S_{ij}} = \sigma_{\text{QP},S}^2 k_{\text{QP}}(t_i, t_j) + \sigma_{\text{WN},S}^2 I_{ij} \quad (2.2)$$

where $\sigma_{\text{QP},S}$ and $\sigma_{\text{WN},S}$ are the two amplitudes, and \mathbf{I} is the identity matrix.

2.2.3 Fitting the S index

With a single, time-independent parameter representing the mean value of the S index, μ_S , the set of fit parameters contains six variables:

$$\boldsymbol{\theta}_S = \{\mu_S, \sigma_{\text{QP},S}, \sigma_{\text{WN},S}, \tau, \eta, P_{\text{rot}}\} \quad (2.3)$$

and the GP likelihood function [72] is given by

$$\mathcal{L}(\boldsymbol{\theta}_S | \mathbf{s}, \mathbf{t}) = \frac{1}{\sqrt{\det(2\pi\mathbf{K}_S)}} \exp\left(-\frac{1}{2}\Delta\mathbf{s}^T \mathbf{K}_S^{-1} \Delta\mathbf{s}\right) \quad (2.4)$$

Fit Parameter	Units	Prior	MCMC Initial Guess	Fit Value
μ_S		$U(0.15, 0.17)$	mean(s) = 0.1626	$0.1627^{+0.0002}_{-0.0002}$
$\sigma_{QP,S}$		$U(10^{-5}, 10^{-1})$	std(s) = 1.83×10^{-3}	$1.43^{+0.10}_{-0.09} \times 10^{-3}$
$\sigma_{WN,S}$		$U(10^{-5}, 10^{-1})$	std(s) = 1.83×10^{-3}	$0.178^{+0.007}_{-0.007} \times 10^{-3}$
τ	days	$U(5, 100)$	22	$23.6^{+1.1}_{-1.1}$
η		$U(0.1, 0.9)$	0.5	$0.58^{+0.04}_{-0.04}$
P_{rot}	days	$U(24, 32)$	27	$28.1^{+0.4}_{-0.5}$
μ_{RV}	m s^{-1}	$U(97, 113)$	mean(\mathbf{r}) = 102.4	$102.4^{+0.3}_{-0.3}$
$\sigma_{\text{QP,RV}}$	m s^{-1}	$U(0.01, 10)$	std(\mathbf{r}) = 1.65	$1.44^{+0.16}_{-0.15}$
$\sigma_{\text{WN,RV}}$	m s^{-1}	$U(0.01, 10)$	std(\mathbf{r}) = 1.65	$1.25^{+0.04}_{-0.04}$

Table 2.1: Summary of parameters used for Markov chain Monte Carlo (MCMC) sampling for the S index and RV fits of solar telescope and HARPS-N data. The upper and lower bounds of the uniform priors are given in addition to the initial guesses used for each parameter. The last column shows the resulting median value of the MCMC samples and their corresponding 16% and 84% quantiles as error bars.

where $\Delta \mathbf{s} = \mathbf{s} - \mu_S$ and \mathbf{s} is a vector containing the daily averaged S index at times \mathbf{t} for all the S index values shown in Fig. 2.1a. Each parameter is assigned a uniform prior probability with upper and lower bounds encompassing realistic values, shown in Table 2.1. The posterior probability distribution, $p_S(\boldsymbol{\theta}_S | \mathbf{s}, \mathbf{t})$, is then proportional to $\mathcal{L}(\boldsymbol{\theta}_S | \mathbf{s}, \mathbf{t})$ within the prior bounds and 0 otherwise.

We estimate the posterior distribution using an affine-invariant Markov chain Monte Carlo (MCMC) method, implemented with the Python [76] packages NumPy [77], SciPy [78], emcee [79], and george [80]. Following [79] 32 walkers are used to sample the parameter space, and are initialized with a normal distribution around mean values from a maximum likelihood fit and known properties of the Sun. These values are summarized in Table 2.1. The first 100 samples are discarded allowing the walkers to converge to the posterior distribution before evaluating an additional 50,000 samples per walker. To remove correlations between samples, they are thinned by keeping one out of every n_S samples for each walker, where $n_S = 21$ is the average correlation length of the walkers as estimated by the autocorrelation. This yields a total of 76,032 uncorrelated samples for each parameter.

The resulting parameter estimates, $\hat{\boldsymbol{\theta}}_S$, are displayed in the last column of Table 2.1 as the median value of these uncorrelated samples. The error bars are reported as the 16% and 84% quantiles. The values obtained are consistent with known solar properties. Of note is the active region lifetime, τ , which is less than one rotation period despite lifetimes of photospheric faculae typically being greater than six rotation periods. Variability of the distribution of faculae on the solar surface is expected to drive this value down below the lifespan of an individual facular region. The GP regression fit to the S index and a histogram of the residuals are shown in Fig. 2.1a,b. The variances and covariances of the MCMC samples are shown graphically in the bottom left half of Fig. 2.2. The largest correlations are the $(\eta, \sigma_{QP,S})$ and $(\eta, \sigma_{WN,S})$ pairs with statistically significant Pearson correlation coefficients [81] of 0.49 and 0.42 respectively. This could indicate an unaccounted for spatial dependence in the magnetic activity. Furthermore, significant correlations exist in the (η, τ) and (η, P_{rot}) pairs, potentially reflecting the migration of active regions from higher latitudes (and thus longer rotation period) to lower latitudes (and thus shorter rotation period) over the 2.2 years of data. Additionally, $\sigma_{QP,S}$ increases slightly with τ , lending further evidence to a time dependence in the fit parameters as τ will be longer for the on average larger, longer-lived active regions earlier

in the solar magnetic cycle. These correlations suggest a more sophisticated model may more closely represent the physical mechanisms in the photosphere of the Sun and thus capture more of the RV variation due to activity. Further investigation of these correlations are left to future study.

2.2.4 Fitting the RVs

We assume that the magnetic activity driving the S index also affects the observed RVs with equivalent correlated variability and thus may be described with equal GP regression parameters: τ , η , and P_{rot} . Unlike [73] who fit the S index and RVs simultaneously, we model the S index and RVs consecutively. We fit the RVs using the same quasiperiodic kernel function of equation (2.1) with the spot lifetime, τ , spot distribution, η , and the rotation period P_{rot} fixed at the median value from the S index fit. Because the effects of solar system planets have already been removed from the RVs, we are left with only three fit parameters:

$$\boldsymbol{\theta}_{\text{RV}} = \{\mu_{\text{RV}}, \sigma_{\text{QP,RV}}, \sigma_{\text{WN,RV}}\} \quad (2.5)$$

and the likelihood function becomes

$$\mathcal{L}(\boldsymbol{\theta}_{\text{RV}}|\mathbf{r}, \mathbf{t}) = \frac{1}{\sqrt{\det(2\pi\mathbf{K}_{\text{RV}})}} \exp\left(-\frac{1}{2}\Delta\mathbf{r}^T\mathbf{K}_{\text{RV}}^{-1}\Delta\mathbf{r}\right) \quad (2.6)$$

where $\Delta\mathbf{r} = \mathbf{r} - \mu_{\text{RV}}$, \mathbf{r} is a vector of the RVs at times \mathbf{t} and \mathbf{K}_{RV} is populated with $\boldsymbol{\theta}_{\text{RV}}$ by substituting all ‘S’ subscripts with ‘R’ in equation (2.2). Priors are shown in Table 2.1.

The posterior probability, $p_{\text{RV}}(\boldsymbol{\theta}_{\text{RV}}|\mathbf{r}, \mathbf{t})$, is proportional to the likelihood function within the bounds of the prior and 0 otherwise. It is sampled using the same MCMC protocol as with the S index and we observe a correlation length of $n_{\text{RV}} = 11$. Retaining one of every n_{RV} samples leads to a total of 145,312 uncorrelated samples. The parameter estimates, $\hat{\boldsymbol{\theta}}_{\text{RV}}$, are summarized in Table 2.1 and the resulting GP regression fit and residuals are shown in Fig. 2.1. The fit reduces the RMS scatter in the data from 1.65 m s^{-1} to 1.14 m s^{-1} , consistent with [82], [83], and other solar analyses (see

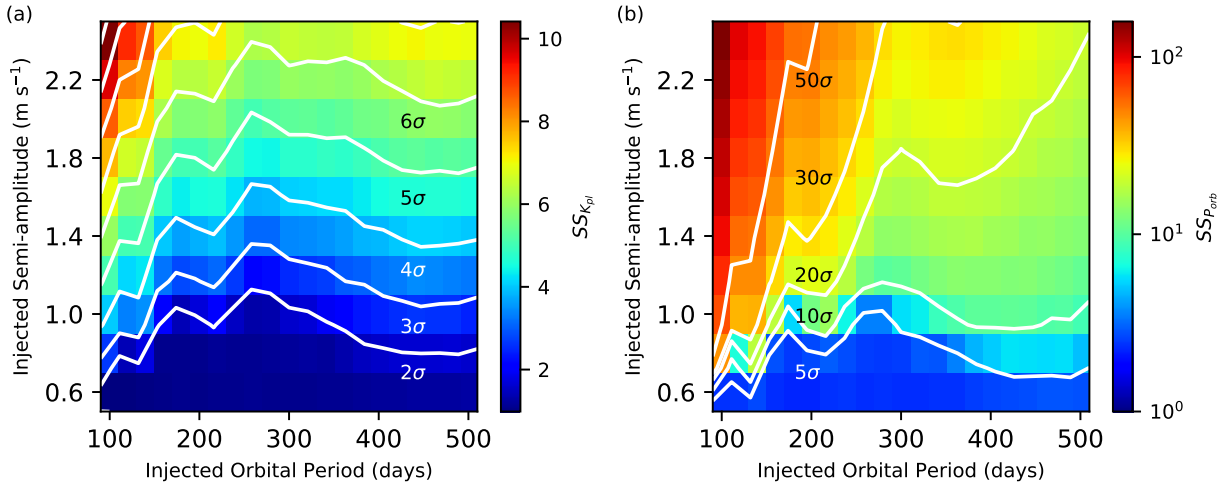


Figure 2.3: Sensitivity maps of the recovered (a) semi-amplitude, K_{pl} , and (b) orbital period, P_{orb} , of synthesized planets using the GP regression with synthetic planetary signals injected into the 800 days of solar RVs. Color bars show the statistical significance (SS) of the recovered parameters as defined in equation (2.11). The white lines in (a) show contours of $2\sigma, 3\sigma, \dots$ statistical significance and the white lines in (b) show contours of $5\sigma, 10\sigma, 20\sigma, 30\sigma,$ and 50σ statistical significance. Each “pixel” in either image represents one of the 210 simulated planets. The orbital period is generally recovered with a high degree of confidence but the semi-amplitude is only recovered at the 5σ level for planets with semi-amplitude above 1 m s^{-1} .

section 2.5). The marginalized distributions are shown in Fig. 2.2 and again display a well behaved posterior distribution. The only significant correlation exists between $\sigma_{\text{QP,RV}}$ and $\sigma_{\text{WN,RV}}$ and is negative, which is expected as these parameters will trade off the amount of variation seen in the RVs.

2.3 Sensitivity map

2.3.1 Synthetic planet model

To explore the limits of our GP regression for detecting low-mass, long-period exoplanets, we inject synthetic planets with varying Keplerian parameters into the solar RVs. The general Doppler induced radial velocity, $v_{\text{RV}}(t)$, of a host star by a companion planet is given by

$$v_{\text{RV}}(t) = K_{\text{pl}} [\cos(\omega + \nu(t)|t_{\text{p}}, P_{\text{orb}}, e)) + e \cos(\omega)] \quad (2.7)$$

where K_{pl} is the semi-amplitude, ω is the argument of periastron, e is the eccentricity, t_{p} is the time of pericenter passage, P_{orb} is the orbital period, and $\nu(t|t_{\text{p}}, P_{\text{orb}}, e)$ is the true anomaly [84]. As a best case scenario, we restrict ourselves to circular orbits (i.e., $e = 0$), which simplifies the Doppler shift to the sine function

$$v_{\text{RV}}(t) = K_{\text{pl}} \sin\left(\frac{2\pi t}{P_{\text{orb}}} + \phi\right) \quad (2.8)$$

where ϕ is an arbitrary phase. A synthetic planet is then generated by choosing values for the semi-amplitude, orbital period, and phase followed by adding $v_{\text{RV}}(\mathbf{t})$ to the vector of measured solar RVs, \mathbf{r} , observed at times \mathbf{t} . The vector of RVs thus undergoes the transformation

$$\mathbf{r} \rightarrow \mathbf{r} + v_{\text{RV}}(\mathbf{t}). \quad (2.9)$$

The GP regression now contains a mean function given by the addition of the overall mean value of the RVs, μ_{RV} , along with the Keplerian parameters required to describe the injected circular planetary orbit:

$$\boldsymbol{\theta}_{\text{RV}} = \{K_{\text{pl}}, P_{\text{orb}}, \phi, \mu_{\text{RV}}, \sigma_{\text{QP,RV}}, \sigma_{\text{WN,RV}}\} \quad (2.10)$$

and the likelihood function is still given by (2.6), with the exception that the vector of fit residuals becomes $\Delta\mathbf{r} = \mathbf{r} - (\mu_{\text{RV}} + v_{\text{RV}}(\mathbf{t}))$ with \mathbf{t} the vector of observation times and v_{RV} the function defined in Eq. 2.8.

2.3.2 Retrieval of injected signals

Using the techniques of the previous section, we construct a map of detection sensitivities for a range of synthetic, low-mass, long-period planets. Our grid contains 210 injected planets with semi-amplitudes from 0.6 m s^{-1} to 2.4 m s^{-1} in 0.2 m s^{-1} steps and orbital periods from 100 days to 500 days in 20 day steps. The phase of each planet is drawn from a uniform distribution, $\phi \sim U(-0.1, 0.1)$ radians, and the prior on the phase is uniform from $-\pi$ to π radians. This is done to avoid numerical instabilities associated with the phase occurring near the boundary of the prior. The prior on the semi-amplitude allows only positive values less than 10 m s^{-1} and the prior on the

orbital period is in the range $(\frac{1}{2}P_{\text{orb}}, 2P_{\text{orb}})$. For planets detected with a high degree of statistical significance, the priors are uninformative and do not affect the results. However, for the lowest mass planets with 1σ detections or less, the priors do constrain the results as described below. These planets set the lower bound on the range of injected semi-amplitudes used in this analysis.

We draw MCMC samples, exploring the semi-amplitude and orbital period linearly, in the same fashion as the previous section with the non-Keplerian parameter priors unchanged. We again define the parameter estimates, $\hat{\theta}_{\text{RV}}$, as the median value of the uncorrelated samples and the corresponding lower and upper bounds, $\hat{\theta}_{\text{R, lower}}$ and $\hat{\theta}_{\text{R, upper}}$, as the 16% and 84% quantiles. The statistical significance, SS_i , is defined as

$$SS_i = \frac{|\hat{\theta}_i|}{\frac{1}{2}(\hat{\theta}_{i, \text{upper}} - \hat{\theta}_{i, \text{lower}})} \quad (2.11)$$

where θ_i is a given parameter from the vector defined in Eq. (2.10). We plot the statistical significance of the recovered orbital period and semi-amplitude of the 210 synthetic planets in Fig. 2.3. The orbital period is determined with a high degree of statistical significance, though for semi-amplitudes below 1 m s^{-1} this should be taken as an upper bound as the MCMC samples begin to encounter the edges of the uniform prior. The semi-amplitude, however, is much less certain. For a 5σ “discovery” threshold, planets with a semi-amplitude less than 1 m s^{-1} would require more observations than the 800 days of solar telescope data used in this analysis. The structure of the contours in Fig. 2.3(a) is likely due to our non-continuous observing schedule. In particular, [85] find that activity-driven signals at orbital periods unrelated to either planetary orbital periods or the stellar rotation period can arise from uneven sampling. Thus even the near-daily, long-baseline observing schedule of the solar telescope decreases the semi-amplitude sensitivity [58]. We find agreement within parameter uncertainties between the injected and extracted orbital periods and semi-amplitudes, suggesting no systematic effects induced by the GP regression. In the next section we explore the baseline of observations required to detect sub- m s^{-1} planets using completely synthetic RVs.

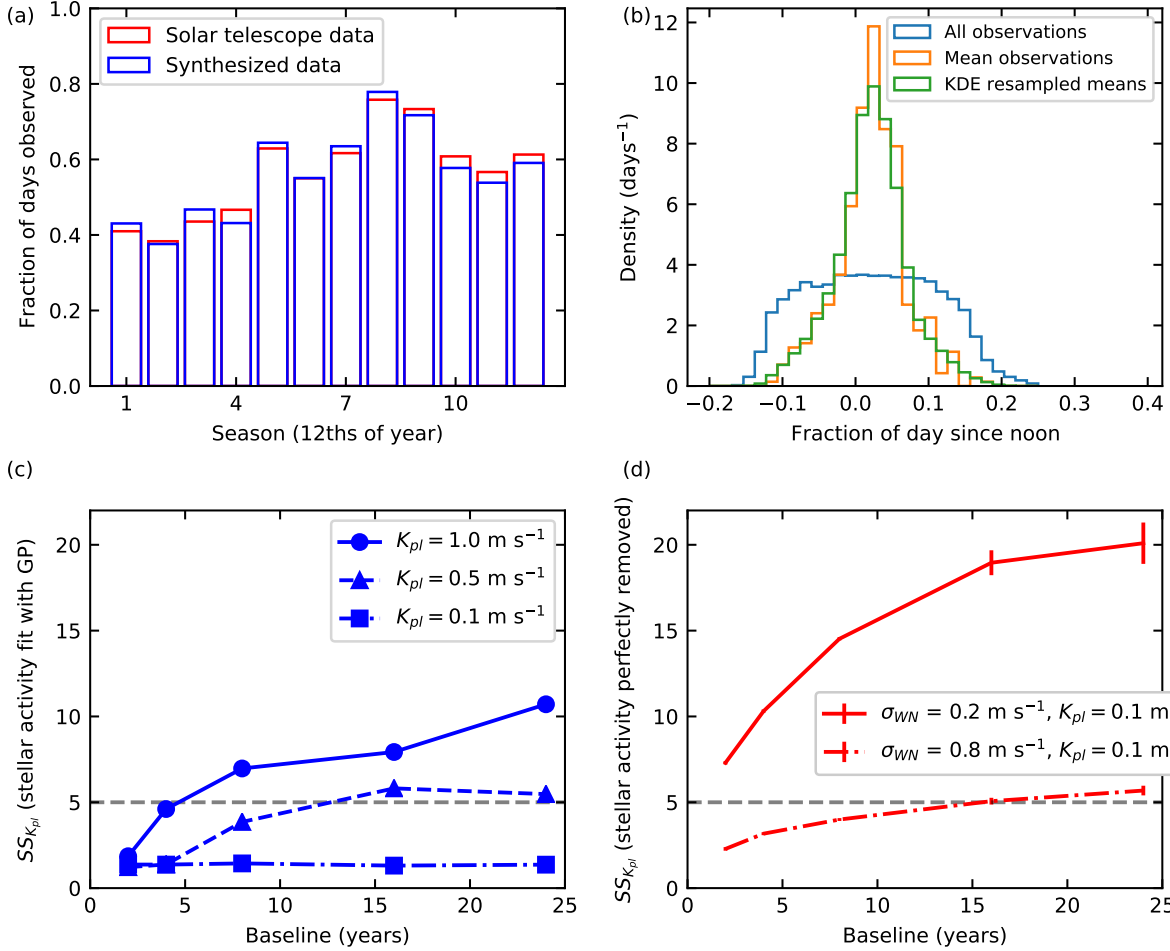


Figure 2.4: Generating synthetic RVs based on historical solar telescope observations. (a) Seasonal variation of solar telescope and synthetic observations. Red histogram shows the fraction of days with at least ten 5 minute exposures for 12 evenly spaced sections of a year (does not correspond to calendar months). Blue histogram shows the analogous ratio of seasonal observations from dates synthesized from the red histogram. (b) Histogram of RV observation times given as a fraction of a day since noon UTC: *blue* - all 5 min solar telescope exposures, *orange* - daily averaged observation times for days with at least ten 5 minute exposures, *green* - synthesized observation times using a Gaussian kernel density estimate. (c) Statistical significance (SS), defined in equation (2.11), of the recovered semi-amplitude, K_{pl} , using a GP regression for an “exo-Venus” (225 day orbital period) as a function of baseline of data. Solid (circles) and dashed (triangles) curves show the SS for an “exo-Venus” with semi-amplitudes of 1.0 m s^{-1} and 0.5 m s^{-1} respectively. Dashed-dotted (squares) curve shows an upper bound for a semi-amplitude of 0.1 m s^{-1} . Dashed grey line shows the 5σ detection threshold. (d) SS for a 0.1 m s^{-1} “exo-Venus” with stellar variability perfectly removed (i.e., only white noise). Solid and dashed curves show the SS for white noise amplitudes of 0.2 m s^{-1} and 0.8 m s^{-1} respectively. Error bars are shown as vertical lines and are too small to be seen for baselines less than 15 years. Again, dashed grey line shows the 5σ detection threshold.

2.4 Synthetic RVs

To study longer observing baselines than in the previous section, we synthesize not only Keplerian Doppler shifts, but also the solar RVs themselves. We thus extend the observing baseline to determine the requirements for detecting a low-mass, long-period analog. A synthetic planet with the orbital period of Venus (225 days) is injected with varying semi-amplitude. We use the orbital period of Venus to avoid systematic biases associated with measuring any periodic signal at 365 days.

2.4.1 Observing schedule

The first step to synthesizing solar telescope data is to create a realistic observing schedule to account for bad weather and telescope downtime. We assume a naive model where each calendar month, m , is assigned an observation probability, $p_{\text{obs}}(m)$, which represents the fraction of days that have observations in month m . These probabilities are estimated using historical solar telescope data with the corresponding fraction of days with 10 or more observations. These probabilities are shown in red in Fig. 2.4a.

For any given day d in month m , a uniform random number, $r(d) \sim U(0, 1)$, is drawn and compared to $p_{\text{obs}}(m)$. If $r(d) \leq p_{\text{obs}}(m)$, an observation occurs. This process is repeated for as many consecutive days as needed. Using this process, we create a synthetic 30 year observing schedule mimicking the solar telescope seasonal variations. The resulting synthetic observing schedule is shown in blue in Fig. 2.4a.

2.4.2 Observation times

The second step for simulating data is to generate an observation time for each day with an observation. We begin by computing the mean observation time of all exposures in each day from

historical solar telescope data, limiting only to days with at least 10 exposures. This distribution of mean observation times is histogrammed and fit to a Gaussian kernel density estimate (KDE) [86]. Synthetic observation times are drawn from this KDE. Fig. 2.4b shows the distribution of mean telescope exposure times and the resulting observation times drawn from the KDE. These observation times coupled with the observing days from the previous section completely determine the observing schedule of the simulated data.

2.4.3 Synthesized RVs

Finally we synthesize the solar telescope RVs using a GP with the same kernel function given in equation (2.2). The parameter values used to populate the covariance matrix are taken from the GP regression fit of the solar data as given in Table 2.1. A random sample of RVs is then drawn from the GP using these parameter values and the synthetic observation times generated in the previous section. Because the solar telescope RVs only cover a fraction of the solar magnetic cycle during solar minimum, we do not model this effect and thus the resulting synthetic RVs represent a best case scenario for this model. Additionally, a Keplerian term given by equation (2.8) is added with several test semi-amplitudes, an orbital period of 225 days, and a phase drawn randomly in the interval $(-0.1, 0.1)$ radians.

We then draw MCMC samples (see appendix A) and define the parameter estimates and lower and upper bounds as before to determine the statistical significance of the recovered semi-amplitude for baselines of data equal to 2, 4, 8, 16, and 24 years. We repeat this process for semi-amplitudes given by 1, 0.5, and 0.1 m s^{-1} . The resulting statistical significances are shown in Fig. 2.4c. The dashed-dotted curve representing the 0.1 m s^{-1} planet represents an upper bound, with the prior probability distribution on the semi-amplitude restricting its value to be positive. Even for a 0.5 m s^{-1} planet, we determine that between 10 and 15 years of data are needed to reach the 5σ discovery threshold using this GP regression and data similar to that of the HARPS-N solar telescope.

As shown above, a multi-decade temporal baseline is required to detect a temperate, low-mass

planet orbiting a Sun-like star using this GP regression. Thus we next assume that we have a direct measure of the magnetic variability of the target star and are able to perfectly remove the effects of variability from the RVs, leaving only the planetary signals and white noise. We can then fit to a sine curve with simple least squares methods. We expect the semi-amplitude uncertainty, $\sigma_{K_{\text{pl}}}$, to scale with the measurement sensitivity, σ_{RV} , divided by the square-root of the number of RV measurements, N [87]:

$$\sigma_K = \sigma_{\text{RV}} \sqrt{\frac{2}{N}}. \quad (2.12)$$

To confirm this scaling, we synthesize RV time series comprising simple white noise and a Keplerian Doppler shift of $K_{\text{pl}} = 10 \text{ cm s}^{-1}$ for an injected planet with an orbital period of 225 days (i.e., an exo-Venus). We generate this data for varying baselines and white noise levels of both 0.8 m s^{-1} and 0.2 m s^{-1} . The statistical significance of the semi-amplitude, K_{pl} , from least squares fits to these data sets are shown in Fig. 2.4d and are in good agreement with expectations set by Eq. 2.12. We emphasize that we need to reduce the white noise to levels approaching 0.2 m s^{-1} to reach the 5σ detection threshold in only a few years for a true exo-Venus.

2.5 Discussion and outlook

The GP regression of Sec. 2.2 reduces the RMS variation of the solar RVs from 1.65 m s^{-1} to 1.14 m s^{-1} . However, treating stellar variability with this GP regression still requires 10 to 15 years of densely sampled RV observations to detect long-period, low-mass planets. This is much longer than would be expected if the RVs contained purely uncorrelated white noise. This result is in line with more physically-motivated techniques. For example, [82] used magnetograms and dopplergrams from the Helioseismic and Magnetic Imager (HMI) on board the Solar Dynamics Observatory (SDO) to derive activity-driven RV time series. By modeling the HARPS-N RVs using these activity time series, they reduced the RV RMS from 1.65 m s^{-1} to 1.21 m s^{-1} . [59] reduced the RV RMS to 0.85 m s^{-1} by modeling the RVs with a linear combination of the unsigned magnetic flux from HMI and the total solar irradiance, using the FF' method [88]. Fitting our GP regression to the unsigned flux may yield a smaller variation in the RV residuals, and this will be investigated in a future work. [89]

and [90] used logistic regression and gradient boosting on HARPS-N solar spectra to differentiate activity sensitive and insensitive lines. Computing RVs from these sets of lines reduced the RV RMS to 0.9 m s^{-1} . Similarly, [83] applied the techniques of [91] to the HARPS-N solar RVs to estimate magnetoconvective RV variations. This analysis, however, was unable to reduce the observed RMS RVs.

That the GP regression — a statistical tool — performs similarly to these more physically-motivated analyses demonstrates its power: rather than requiring high-resolution solar images or specially tuned line lists, the GP regression assumes a correlation between the solar S index and observed RVs. However, the fact that the GP analysis and these physically-motivated techniques arrive at the same approximately 1 m s^{-1} RV uncertainty level indicates that something else — either another physical process operating on a different timescale or an instrumental systematic — is limiting the performance of these techniques. [59] discuss some physical effects missing from the current models, to be addressed in future work. However for current data analysis models, the solar dataset used in this work, with nearly 2.5 years of near-daily observations mostly in the decline phase of Cycle 24, represents a best-case scenario; we can therefore only perform 5σ recoveries of long-period planets with semi-amplitudes greater than 1 m s^{-1} , consistent with the results of the recent community-wide RV challenge of [60].

Our work using synthetic RVs indicates that more observations will not quickly overcome this limit: As shown in Fig. 2.4, it will take 10-15 years to get 5σ on a 0.5 m s^{-1} RV signal with a 225 day period (i.e., the orbital period of Venus), and at least 25 years for a 0.1 m s^{-1} RV signal. The last panel of this figure also indicates that a *perfect* model of activity-driven correlated variations would not in itself suffice for the rapid detection of an exo-Venus or an exo-Earth: even in the absence of correlated noise, a current-generation spectrograph with a long term stability of about 0.8 m s^{-1} would need a 10-15 year observing baseline to reach a 5σ detection of an exo-Earth. Successful exo-Earth discovery therefore requires both more sophisticated models of stellar variability and the improved RV precision, long term stability, and dense observational sampling from a next-generation spectrograph [92] such as ESPRESSO [93], NEID [94], EXPRES [95], HARPS3 [96], or G-CLEF [97] on the GMT.

In addition to applications in exoplanet searches, GPs are employed in many areas of astrophysics. Notably the same GP model discussed in this chapter is employed in a proposal to map the local galactic dark matter density by measuring the RV of a cluster of stars and comparing the radial galactic acceleration to that of the Sun. Differences in acceleration that deviate from a model excluding dark matter can help probe the distribution of dark matter in the Milky Way. Details are provided in [appendix B](#).

Solar atmospheric differential extinction corrections

3.1 Introduction

Over the course of a day, the sun will experience systematic RV shifts of about ~ 1 m/s due to the top and bottom of the solar disk experiencing different paths through the atmosphere. Furthermore, because the top half and bottom half of the solar disk will have different redshifts due to the rotation of the sun, we also expect to see a wavelength dependent differential atmospheric extinction. We anticipate this effect could be as much as another 10 cm/s on top of the 1 m/s. This effect is thus relevant to our ongoing search for noise sources in the solar RV data.

The amount of atmosphere a photon traverses is called airmass, labeled A . We define $A = 1$ as the amount of atmosphere traversed from sea level straight up through the zenith to infinity. Thus in general, an observer at sea level will see a celestial body with an airmass $A \geq 1$. Note this means an observer above sea level could observe a celestial body with $A < 1$.

As the photons traverse the atmosphere, they will be scattered and absorbed by the atmosphere. It follows that a given object will be dimmer at greater airmass. An object will experience a dimming

given by $\Delta m = m_f - m_i = k(\lambda)A$ where m_i is the magnitude of the object in space above the atmosphere, m_f is the observed magnitude, and $k(\lambda)$ is the extinction coefficient in magnitudes per airmass. The extinction coefficient is a parameter that depends on the properties of the atmosphere and the wavelength of the light.

We can see now the origin of atmospheric extinction. The top half of the solar disk will experience a smaller airmass than the bottom half and thus will appear brighter. Unless the rotation axis of the sun is aligned perfectly in the plane of the sky (which in general it will not), the rotation of the sun will lead to a differential RV shift between the top and bottom halves of the solar disk. Given normal observing conditions, this can lead to an RV shift up to ~ 1 m/s, which will need to be corrected for.

The current correction procedure assumes $k(\lambda) = k_0$. In other words, the wavelength dependence of the extinction coefficient is not accounted for. The goal here is to study the wavelength dependence of the extinction coefficient to determine the size of this effect and whether or not it needs to be corrected for. If the RV shift is large enough, we can modify the current extinction correction procedure to include a correction for this effect.

[98] shows that for observations in the optical and nearby wavelengths, there are three main contributions to the wavelength dependence of the extinction coefficient:

1. Rayleigh scattering: this is well understood and modeled. ($1/\lambda^4$ dependence)
2. Absorption from Ozone and Water: water is mostly irrelevant for the HARPS-N observation band, but there is a significant ozone absorption band centered around 600 nm.
3. Aerosol scattering: this includes scattering off dust, salt, etc. This is the trickiest one to model and varies wildly from night to night.

Ch. 6, Sc. 56 of [99] gives an estimate of each of each effect at various wavelengths. The result is summarized in Fig. 3.1:

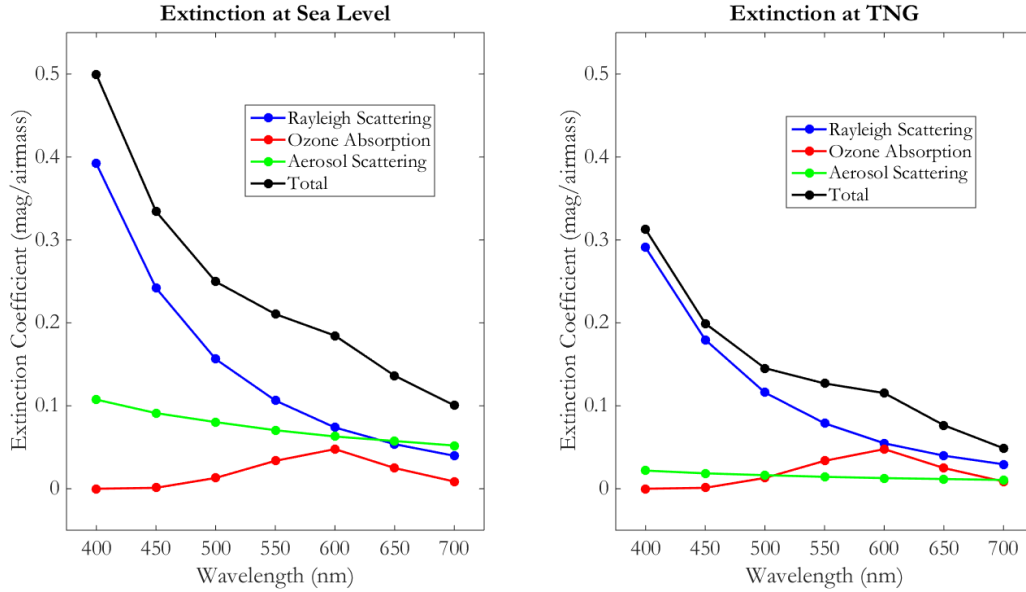


Figure 3.1: Wavelength dependent extinction coefficient in the visible for various sources at sea level (left) and at the same altitude as the TNG (right).

The Rayleigh and Aerosol scattering terms are reduced at the TNG because this site is almost 3km above sea level and thus the photons traverse a thinner atmosphere on average. The Ozone absorption remains unchanged because the majority of the ozone is contained between 10 and 35 km above sea level, far above the TNG. This gives us an idea of what we should expect to see when we measure the extinction coefficient at various wavelengths.

3.2 Derivation of differential extinction

3.2.1 Limb darkening

Let us begin assuming no atmospheric extinction. An effect called limb darkening causes the light from the solar disk to appear dimmer as you move radially outward toward the edge (or limb) of the disk. This arises from the fact that photons near the limb of the disk are emitted higher in the photosphere of the sun where the temperature is cooler.




Figure 3.2: Limb darkening diagram. Earth is in the direction of the arrow point away from A. Credit: [wikipedia](#)

If the mean free path of photons in the photosphere is L , then on average photons that reach the earth will have been emitted from this depth in the photosphere. Shown in Fig. 3.2, photons near the edge of the disk as viewed from Earth must have been emitted higher in the photosphere. It is cooler in the upper photosphere and thus the intensity will decrease due to blackbody radiation.

There are several models to describe this effect, but due to the complicated nature of the solar photosphere, astronomers usually use a one parameter model that works sufficiently well:

$$I(r)/I_0 = 1 - u + u\sqrt{1 - r^2} \quad (3.1)$$

where u is the limb darkening coefficient and r is the radial distance of the 2D solar disk as viewed by an observer on Earth in units of the radius of the solar disk. For most purposes, this model with a coefficient of $u \approx 0.6$ works well.




Figure 3.3: Rotation of the (x, y, z) frame to the (u, v, w) frame about the x -axis.

We may now incorporate this into our model of RV shifts in the solar data. We begin by placing the 2D solar disk in the $x - y$ plane such that the rotation axis of the sun is along \hat{y} , the observer is along \hat{z} , and the origin is at the center of the disk. Units will be such that 1 is equal to the radius of the disk. We can obtain the RV shift by integrating over the disk:

$$RV_{\text{shift}} = \frac{1}{N} \int_{-1}^1 \int_{-\sqrt{1-y^2}}^{\sqrt{1-y^2}} v_z(x, y) (1 - u + u\sqrt{1-r^2}) dx dy \quad (3.2)$$

where $N = \pi(1 - u/3)$ is a normalization constant obtained by integrated the limb darkening term over the disk, and $v_z(x, y)$ is the z -component of the surface velocity of the sun at the point (x, y) .

3.2.2 Radial velocity term

We now rotate the coordinate axes about the x axis by an angle θ , shown in Fig. 3.3. Moving back and forth between the two frames will allow us to take into account the fact the north pole of the Sun will not always be at the edge of the 2D disk. First, if the rotation matrix that describes the rotation from the new (u, v, w) frame to the old (x, y, z) frame is $R(\theta)$ then we have

$$\vec{v}_{old} = R(\theta)\vec{v}_{new} \Rightarrow v_{z,old} = \cos(\theta)v_{w,new} = d_{SNP}v_{w,new} \quad (3.3)$$

where d_{SNP} is the distance from the x axis to the solar north pole in the 2D disk in units of radius. Now we just need to figure out $v_{w,new}$. Let us first figure out the magnitude of the velocity. A point on the surface of the Sun will be traveling in a circle with a radius given by the projection of its position vector onto the $u - w$ plane. Thus we have

$$|\vec{v}|^2 = \omega_{Sun}^2 r^2 = \omega_{Sun}^2 (u^2 + w^2) \quad (3.4)$$

where ω_{Sun} is the angular velocity of the Sun. Secondly we know that in the (u, v, w) frame, the velocity will have no v_v component as this is the axis of rotation. This gives us another condition:

$$|\vec{v}|^2 = v_x^2 + v_w^2 \quad (3.5)$$

Next we know the norm of the of the position vector is 1:

$$u^2 + v^2 + w^2 = 1 \quad (3.6)$$

And finally we know the dot product of the position vector and the velocity vector will be 0:

$$uv_u + wv_w = 0 \quad (3.7)$$

This gives us a system of equations we can solve which gives $v_{w,new} = u\omega_{Sun}$. Fortunately, since our rotation was about the x axis, we know $x = u$. This gives us the final result

$$v_z(x, y) = xd_{SNP}R_{eq}\omega_{Sun} \quad (3.8)$$

where we have converted from our units of solar radius to distance by multiplying by the equatorial radius of the Sun R_{eq} .

3.2.3 Atmospheric extinction

Photons passing through the atmosphere will experience a change in magnitude linear in the airmass A with a proportionality constant given by the extinction coefficient $k(\lambda)$. This can be expressed in terms of intensity as

$$I(A)/I_0 = e^{-(\ln(10)/2.5)k(\lambda)A} \quad (3.9)$$

The constant in front of the extinction coefficient arises from the definition of magnitudes in astronomy. Magnitude is defined such that a change in magnitude of 5 is a change in intensity of 2 orders of magnitude. We can get an approximation of the change in intensity from the top to the bottom of the solar disk by taking the derivative:

$$dI(A)/I(A) = -\frac{\ln(10)}{2.5}k(\lambda)dA \quad (3.10)$$

Now we need to compute the change in airmass across the solar disk dA . The simplest approximation for airmass is the plane-parallel approximation shown graphically in Fig. 3.4.

Here we assume the sky is a flat plane above the observer. The airmass is the ratio of the distance the photon travels through to atmosphere to the height of the atmosphere. The angle between the

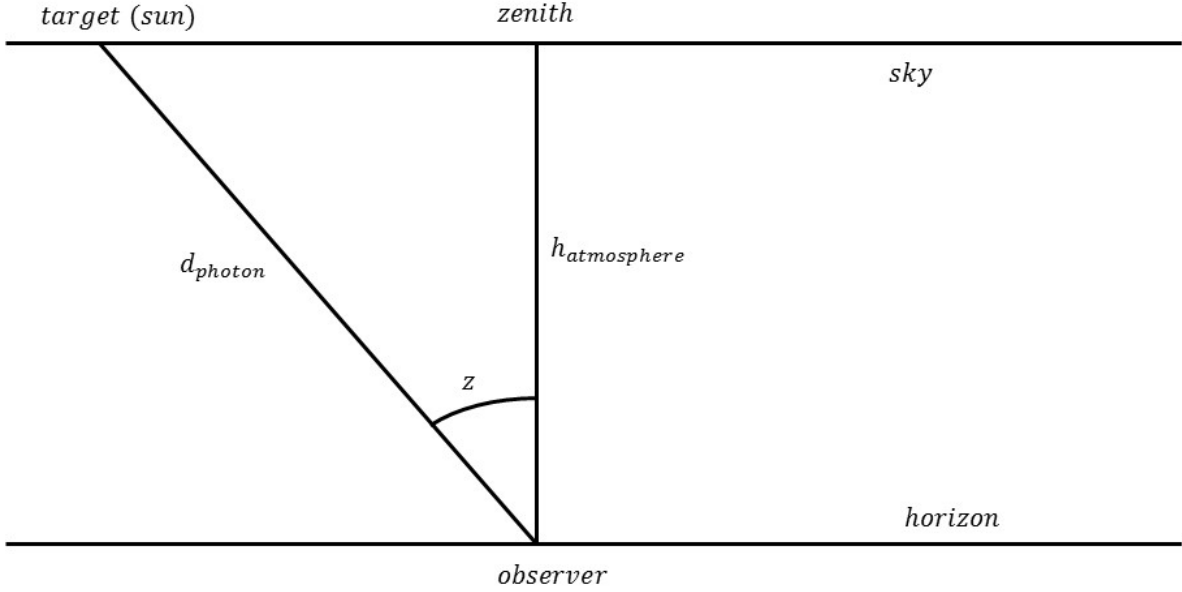


Figure 3.4: Plane parallel approximation for computing airmass.

two is known as the zenith distance and is labeled z in Fig. 3.4. In this approximation it is easy to see the airmass is given by $A = \sec(z)$. This approximation breaks down for fairly small zenith distances though and thus we will adopt the model by [100]:

$$A = (\cos(z) + 0.025e^{-11 \cos(z)})^{-1} \quad (3.11)$$

This approximation provides a good estimate for $z < \pi/2$, more than sufficient for solar RV data.

Now we need to incorporate this into our equation for the RV shift. In general the north pole of the Sun will not align with the great circle connecting the center of the solar disk to the zenith, shown in Fig. 3.5. We can apply a constant slope change in intensity along the direction of changing airmass given by

$$-dI(A)/2 = \frac{\ln(10)}{5} k(\lambda) dA \quad (3.12)$$




Figure 3.5: Diagram showing the misalignment of the rotation of the Sun with the great circle connecting the horizon and zenith.

We will thus get an RV shift given by

$$RV_{\text{shift}} = (d_{SNP} R_{eq} \omega_{Sun} / N) \int_{-1}^1 \int_{-\sqrt{1-y^2}}^{\sqrt{1-y^2}} f(x, y | u, \lambda, dA, \phi) dx dy \quad (3.13)$$

where

$$f(x, y | u, \lambda, dA, \phi) = x [1 - u + u \sqrt{1 - x^2 - y^2}] \left[1 + \frac{\ln(10)}{5} k(\lambda) dA (-x \sin(\phi) + y \cos(\phi)) \right] \quad (3.14)$$

and N is the same normalization constant defined in Eg. 3.2. Evaluating this integral gives the final result:

$$RV_{\text{shift}} = \ln(10)d_{SNP}R_{eq}\omega_{Sun}k(\lambda)dA(15 - 7u) \sin(\phi)/[100(u - 3)] \quad (3.15)$$

3.2.4 Calculation of the angle ϕ

To start off it is important to know that the positions of astronomical objects are usually described in Equatorial coordinates, given by Right Ascension and Declination $(RA, dec) = (\alpha, \delta)$. Imagine extending the Earth's equator and north pole out into the celestial sphere. RA and Dec are then analogous to longitude and latitude respectively. The [JPL Horizons Ephemeris Web Interface](#) provides the RA and Dec of any celestial body in the solar system at a given date and time from any position on Earth.

But here, what we care about is the altitude, a , of the sun, which is the angle the center of the sun and the horizon make with the observer. We can use a transformation of coordinates:

$$\sin(a) = \sin(\delta) \sin(\phi_0) + \cos(\delta) \cos(\phi_0) \cos(\text{GMST} - \lambda_0 - \alpha) \quad (3.16)$$

where ϕ_0 and λ_0 are the observer's latitude and longitude respectively. GMST is Greenwich Mean Sidereal Time. It defined by the international astronomical union as:

$$\text{GMST} = 24110.54841 + 8640184.812866T + 0.093104T^2 - 0.0000062T^3 \quad (3.17)$$

where T is the number of Julian centuries since noon on Jan. 1, 2000:

$$T = (\text{JD} - 2451545.0)/36525 \quad (3.18)$$

JD is the Julian Date (i.e., the number of days since noon on Jan. 1, 4713 B.C.). Now we can compute the parallactic angle, η , which is the angle the zenith and the north celestial pole make with the center of the sun. This is given by

$$\sin(\eta) = \sin(\text{GMST} - \lambda_0 - \alpha) \cos(\phi_0) / \cos(a) \quad (3.19)$$

And finally the angle of interest is just

$$\phi = \eta - \text{NPA} \quad (3.20)$$

where NPA is the north pole angle, another quantity that may be queried with the [JPL Horizons Ephemeris Web Interface](#). This is the angle the solar north pole and the celestial north pole make with the center of the solar disk.

Thus if we are given a time of observation, JD, and the position of the observer on Earth, (ϕ_0, λ_0) , we can query Horizons for the (α, δ) coordinates of the sun and the NPA. We then have a systematic way of computing ϕ . We can also query Horizons for the quantity d_{SNP} , the distance of the solar north pole to the center of the sun on the 2D disk in units of solar radius. We now have everything we need to calculate the RV bias correction in Eq. 3.15.

3.3 Fitting the RV data

If we fit the solar telescope RV data to Eq. 3.15 we get the result shown in Fig. 3.6. We can say the fit agrees with the expected result, but the 95% confidence interval is large so it seems bold to make such a statement. The fit does agree well in the region where we have an abundance of SNR ($500 \text{ nm} < \lambda < 600 \text{ nm}$). But for $\lambda < 500 \text{ nm}$ we have poor agreement as we have very little SNR in this




Figure 3.6: Fit of the wavelength dependent extinction coefficient to the solar telescope RV data. The expected value is shown in red while the fit is shown in blue. The blue shaded region shows the 95% confidence interval.

regime. There are additional problems with this fit, however. When we investigate the residuals, several potential problems arise. First, the residuals are far from Gaussian as shown in Fig. 3.7.

Another problem with this fit is that the result is nonphysical. Fig. 3.8 shows the individual components of the wavelength dependent extinction coefficient fit. The Rayleigh scattering term is consistent with zero and should be the dominant term. Furthermore, the ozone term is less than zero which would imply the light gets more intense as it passes through the Earth’s atmosphere. For these reasons, the fit cannot be trusted. We believe the lack of SNR in the short wavelength regime is responsible for these problems.

3.4 Daily extinction correction used in the RV data

Because the wavelength dependent extinction cannot be corrected for, we correct the RV data with a wavelength independent extinction which we derive here with additional terms not covered in section 3.2.

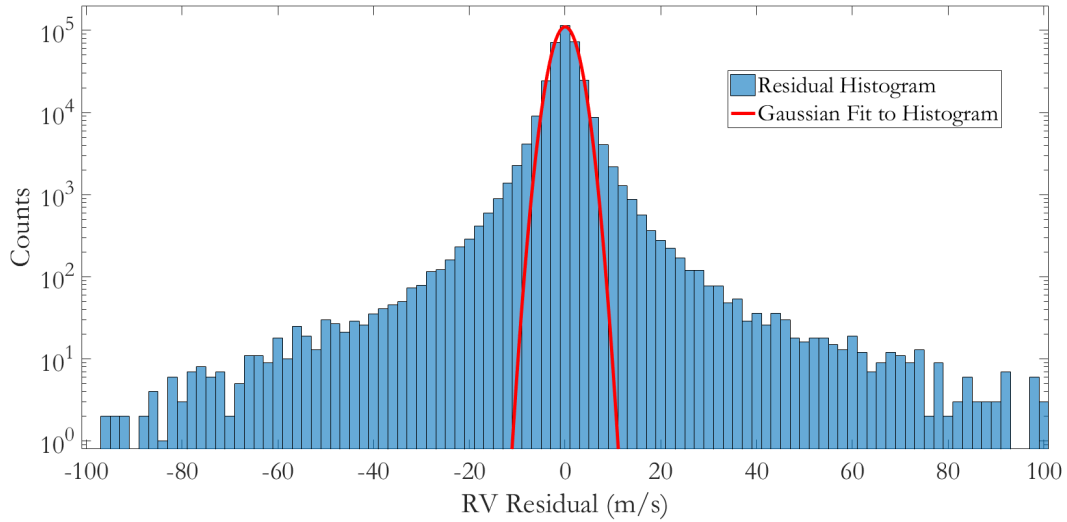


Figure 3.7: RV residuals for the wavelength dependent extinction coefficient fit. The residuals are non Gaussian as shown by the red curve.




Figure 3.8: Individual components of the extinction coefficient fit. The Rayleigh scattering term is shown in blue and is consistent with zero. The ozone term is shown in red and less than zero. The dust scattering term is shown in green.

The daily mean extinction coefficient defines the vertical extinction gradient across the diameter of the solar disk as a function of airmass. This gradient affects the measured solar radial velocity (RV) whenever the solar rotation axis is inclined relative to the vertical direction. This generally leads to a spurious redshift in the morning, when the approaching hemisphere of the Sun is lower in the sky than the receding hemisphere. The effect was first identified and corrected by [101], and has subsequently been modeled by the BiSON team [102]. The analytic approach we adopt here was developed independently of, and is slightly different from, these previous studies, but the results obtained are essentially identical.

The difference in airmass X between the upper and lower limbs of the Sun is

$$\delta X = X_{\text{upper}} - X_{\text{lower}} \quad (3.21)$$

where $X = \sec(z)[1 - 0.0012(\sec^2(z) - 1)]$ [103]. The zenith angles $z_{\text{lower,upper}} = z_{\odot} \pm \sin^{-1}(R_{\odot}/d_{\odot})$ depend on the zenith angle z_{\odot} of the solar center, the solar radius R_{\odot} and the Earth-Sun distance d_{\odot} at the time of observation.

The corresponding difference in magnitudes of atmospheric extinction is $\delta m_{60} = k_{60}\delta X/2$ across the solar radius. We define a fractional difference in brightness $\epsilon \simeq \delta m_{60} \ln(10)/2.5$, with $\epsilon = 0$ at the center of the solar disc.

The angle between the solar rotation pole and the local vertical is $\phi = PA - q$, where PA is the position angle of the solar rotation axis relative to the north celestial pole at the time of observation. The parallactic angle q is the angle subtended at the Sun by the celestial pole and the zenith, giving

$$\sin(q) = \sin(H) \cos(\beta) / \sin(z_{\odot}) \quad (3.22)$$

$$\cos(q) = \sqrt{1 - \sin^2(q)}, \quad (3.23)$$

where H is the solar hour angle and β is the latitude of the observer.

For a linear limb-darkening law with limb darkening coefficient u , the form of the rotation profile in the absence of differential extinction is [104]:

$$f(x) = \int_{-\sqrt{1-x^2}}^{\sqrt{1-x^2}} (1 - u + u\mu(x, y)) dy, \quad (3.24)$$

where the brightness is $I(x, y) = 1 - u + u\mu(x, y)$ and the foreshortening cosine $\mu(x, y) \equiv \sqrt{1 - x^2 - y^2}$.

The orthogonal Cartesian coordinates x and y are in the directions of increasing solar rotation velocity and the solar north pole respectively in the plane of the sky. For solid-body rotation with equatorial velocity v_{eq} , the brightness-weighted correction to the solar RV is

$$\delta v_r = v_{\text{eq}} \frac{\int_{-1}^1 x f(x) dx}{\int_{-1}^1 f(x) dx}. \quad (3.25)$$

In the absence of differential extinction across the rotation profile, the integrand in the numerator of Eq. 3.25 is an odd function, so the brightness-weighted velocity offset is zero.

We modify the calculation of the limb-darkened rotation profile to include such a density gradient, parameterized by the fractional extinction difference ϵ across the solar radius. A brightness gradient of uniform slope ϵ in x and y at an angle ϕ to the local vertical modifies the rotation profile in Eq. 3.25 to a new form:

$$g(x, \phi) = \int_{-\sqrt{1-x^2}}^{\sqrt{1-x^2}} I(x, y) [1 + \epsilon(-x \sin(\phi) + y \cos(\phi))] dy. \quad (3.26)$$

The sign of the $x \sin(\phi)$ term should be negative if ϕ is positive west of the meridian.

The velocity bias δv_r introduced by the extinction gradient thus depends on v_{eq} , u , ϕ , and ϵ according to

$$\begin{aligned} \delta v_r &= \frac{\int_{-1}^1 x g(x, \phi) dx}{\int_{-1}^1 g(x, \phi) dx} v_{\text{eq}} \\ &= \frac{(7u - 15)\epsilon \sin(\phi)}{20(3 - u)} v_{\text{eq}}. \end{aligned} \quad (3.27)$$

So far we have assumed that the solar photosphere rotates as a solid body, ignoring the effects of differential rotation, axial inclination and centre-to-limb variations in convective blueshift. [105] modeled the effects of these modifications to the solar rotation profile on the solar RV in the presence of an extinction gradient. Following their example, we can include differential rotation in our model by allowing the apparent rotational angular frequency ω_{obs} to vary as a function of heliographic latitude b . This alters Eq. 3.25 such that the intrinsic solar equatorial rotation speed v_{eq} is replaced by $R_{\odot}\omega_{\text{rot}}(x, y)$. We adopt the sidereal differential rotation law of [106] as a function of heliographic latitude b , defining

$$\omega_{\text{rot}}(\sin(b)) = A + B \sin^2(b) + C \sin^4(b), \quad (3.28)$$

with $A = 2.972 \times 10^{-6} \text{ rad s}^{-1}$, $B = -0.484 \times 10^{-6} \text{ rad s}^{-1}$ and $C = -0.361 \times 10^{-6} \text{ rad s}^{-1}$. When the solar rotation axis is inclined at angle i to the line of sight, a point at location (x, y) on the visible solar hemisphere has

$$\sin(b) = y \sin(i) + \sqrt{1 - x^2 - y^2} \cos(i). \quad (3.29)$$

The integration for v_{rot} becomes

$$\delta v_{\text{rot}} = \frac{\int_{-1}^1 x h(x, \phi) dx}{\int_{-1}^1 g(x, \phi) dx} R_{\odot} \sin(i) \quad (3.30)$$

where

$$h(x, \phi) = \int_{-\sqrt{1-x^2}}^{\sqrt{1-x^2}} I(x, y) \omega_{\text{rot}}(x, y) [1 + \epsilon(-x \sin(\phi) + y \cos(\phi))] dy. \quad (3.31)$$

To take the Earth's instantaneous orbital motion and the centre-to-limb variation in convective RV into account, we must include additional terms:

$$\delta v_{\text{r}} = \frac{\int_{-1}^1 (x h(x, \phi) R_{\odot} \sin(i) + c(x, \phi)) dx}{\int_{-1}^1 g(x, \phi) dx} - \frac{\int_{-1}^1 x g(x, \psi) dx}{\int_{-1}^1 g(x, \psi) dx} R_{\odot} \omega_{\text{orb}}. \quad (3.32)$$

Here ψ is the position angle between ecliptic north and the local vertical at the Sun's location. It always differs from ϕ by less than the 7-degree tilt of the solar rotation axis relative to the ecliptic.

The position angle of ecliptic north relative to celestial north at the Sun is given by

$$\sin \text{PA}_{\text{ecl}} = \sin \varepsilon \cos \alpha \quad (3.33)$$

where ε is the 23.4-degree obliquity of the ecliptic and α is the solar RA at the time of observation. By analogy with the position angle of the solar rotation pole, $\psi = \text{PA}_{\text{ecl}} - q$ where q is the parallactic angle. The integration is specified separately in Eq. 3.32 in a Cartesian coordinate system in which y is oriented toward the ecliptic north pole and x is in the direction of Earth's orbital motion. The solution has the same form as Eq. 3.27 with v_{eq} replaced by $-R_{\odot} \omega_{\text{orb}}$.

The convective contribution is

$$c(x, \phi) = \int_{-\sqrt{1-x^2}}^{\sqrt{1-x^2}} I(x, y) v_{\text{conv}}(\mu(x, y)) [1 + \epsilon(-x \sin(\phi) + y \cos(\phi))] dy. \quad (3.34)$$

Both $I(x, y)$ and $v_{\text{conv}}(\mu(x, y))$ are azimuthally-symmetric functions of radius only. The contribution of $c(x, \phi)$ to δv_r in Eq. 3.32 is therefore independent of ϵ and ϕ , being a constant that depends on the radial variations in limb darkening and convective velocity only. We therefore treat the effect of the centre-to-limb variation in convective velocity as a time-invariant offset to the velocity zero-point, and evaluate the integrals in Eqs. 3.30 and 3.32 to obtain

$$\begin{aligned} \delta v_r &= \epsilon \sin(\phi) R_{\odot} \sin(i) \left(\frac{A(7u - 15)}{20(3 - u)} \right. \\ &+ \frac{B(19u - 35 + \cos^2(i)(3u - 35))}{280(3 - u)} \\ &+ \left. \frac{C(187u - 315 - \cos^2(i)(630 - 118u) - 105 \cos^4(i)(u - 1))}{6720(3 - u)} \right) \\ &- \frac{(7u - 15)\epsilon \sin(\psi)}{20(3 - u)} R_{\odot} \omega_{\text{orb}} \end{aligned} \quad (3.35)$$

The velocity correction for differential extinction varies almost linearly with airmass, and has an amplitude of 2 to 3 m s^{-1} in an uninterrupted day's observations. As $\cos(i)$ varies through the year, the values of δv_r depart by only 1 or 2 mm s^{-1} from the solid-body rotation values obtained using a sidereal rotation rate $\omega_{\text{sid}} = 2.875 \times 10^{-6} \text{ rad s}^{-1}$. The inclination dependence involves only even

powers of $\cos(i)$, and does not change sign when the Earth crosses the solar equator.

Magnetic field fingerprinting of integrated circuit activity with a quantum diamond microscope

4.1 Introduction

Securing integrated circuits against manufacturing flaws, hardware attacks, and software attacks is of vital importance to the semiconductor industry [107]. Hardware attacks often modify the physical layout of an integrated circuit, thereby changing its function. This type of attack can occur at any stage of the globalized semiconductor supply chain, and can range from insertion of malicious Trojan circuitry during the design and fabrication stages [108], to modification or counterfeiting during packaging and distribution stages [109]. Horizontal integration of the industry has led to contracting of integrated circuit fabrication, packaging, and testing to offshore facilities, resulting in a reduction of secure oversight and quality control [110]. Additional growth of the secondhand electronics market has led to a drastic increase in counterfeit integrated circuits [111]. Detection of integrated circuit tampering or counterfeiting has consequently become essential to ensure hardware can be trusted. Similar issues affect quality control of unintended manufacturing flaws.

Magnetic field emanations from integrated circuits afford a powerful means for non-destructive phys-

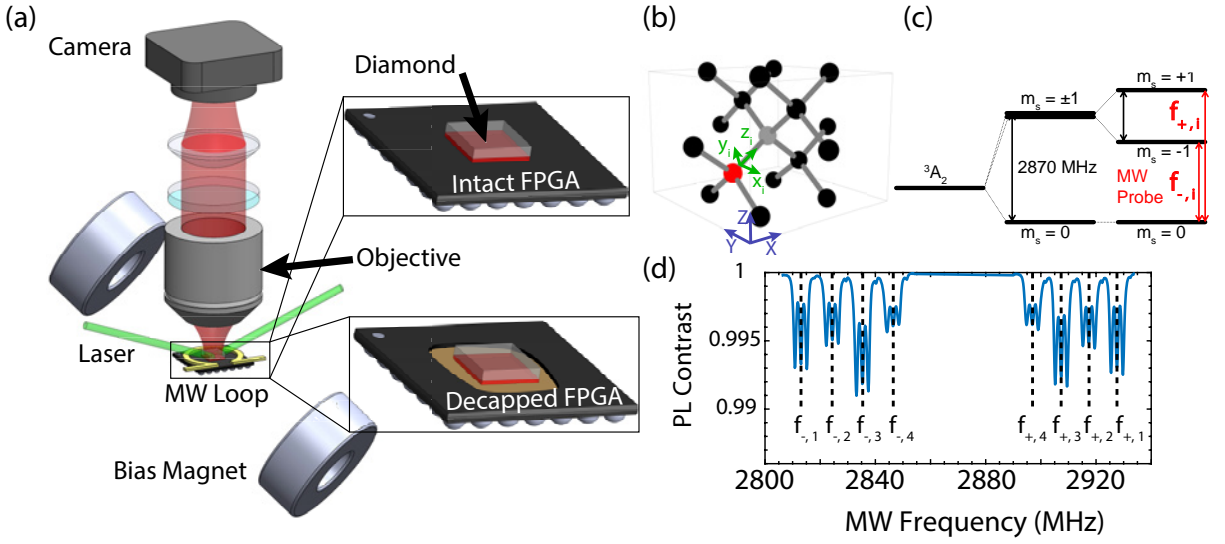


Figure 4.1: (a) Schematic of the Quantum Diamond Microscope (QDM) experimental setup with insets showing the diamond in contact with intact and decapsulated FPGAs. The diamond is positioned such that the NV layer is in direct contact with the FPGA, as indicated by the red layer in the insets. (b) Diamond crystal lattice with the nitrogen (red) vacancy (grey) defect. Lab frame coordinates (X, Y, Z) and NV frame coordinates for a single defect (x, y, z) are shown. (c) The ground state energy level diagram for an NV with fine structure and Zeeman splitting. (d) Example ODMR spectral data for an applied bias field of $(B_X, B_Y, B_Z) = (2.0, 1.6, 0.7)$ mT, showing resonance frequencies of $f_{\pm, i}$ with $i = 1, 2, 3, 4$ indicating each of the four NV axes. Hyperfine interactions between the NV^- electrons and the spin-1 ^{14}N nucleus result in splitting of each NV resonance into three lines.

ical testing. Magnetic fields are generated by current densities in integrated circuits resulting from power and clock distribution networks, input/output lines, word and bit lines, and switching transistors. These currents are present in all operating logic and memory chips and can be leveraged for studying the operational behavior of an integrated circuit during task execution. In general, the resulting integrated circuit magnetic fields pass through many standard integrated circuit materials, and will vary spatially and temporally in ways that correlate with both integrated circuit architecture and operational state. Thus, combined high-resolution and wide-field-of-view mapping of magnetic fields may yield simultaneous structural and functional information, and may be suitable for identification of malicious circuitry or Trojans [112, 113], counterfeit detection [114], fault detection [115–117], and manufacturing flaws [118]. However, leveraging magnetic field emanations is challenging due to the tremendous complexity of circuits integrating billions of transistors of minimum feature sizes down to tens of nanometers, with interconnects distributed across multiple levels of metallization [119]. Multi-layered metal interconnects and three-dimensional stacking give rise to complex magnetic field patterns that are difficult to invert; and large stand-off distances of magnetometers reduce amplitudes of magnetic fields and spatial resolution [120].

We demonstrate how these challenges can be approached using a Quantum Diamond Microscope (QDM) [121–123] augmented with machine-learning classification techniques. With the QDM, we perform simultaneous wide field-of-view, high spatial resolution, vector magnetic field imaging of an operational field-programmable gate array (FPGA). FPGAs are configurable integrated circuits that are commonly used for diverse electronics applications. Systematic and controlled variation of the circuit activity in the FPGA generates complex magnetic field patterns, which we image with the QDM. The QDM employs a dense surface layer of fluorescent nitrogen-vacancy (NV) quantum defects in a macroscopic diamond substrate placed on the integrated circuit under ambient conditions. The electronic spins associated with NV defects have well-established sensitivity to magnetic fields [124–126].

We use the QDM to image magnetic fields from both decapsulated (decapped) and through-package (intact) FPGAs under operational conditions using continuous wave (CW) optically detected magnetic resonance (ODMR) NV spectroscopy. For the decapped FPGA, our measurements yield

magnetic field maps that are distinguishable between operational states over approximately a $4 \text{ mm} \times 4 \text{ mm}$ field-of-view with a 20 nT noise floor, and $\sim 10 \text{ }\mu\text{m}$ magnetic field spatial resolution, limited by the thickness of the NV surface layer in the diamond and the distance to the nearest metal layer. For the intact FPGA, the QDM measurements provide magnetic field maps with a similar field-of-view, 2 nT noise floor, and $\sim 500 \text{ }\mu\text{m}$ magnetic field spatial resolution, limited by the stand-off distance between the NV layer and the FPGA current sources. In particular, we find that operational states of the intact FPGA are distinguishable in the QDM images, even with the diminished magnetic field amplitude and spatial resolution that arise from the large stand-off between the diamond and the integrated circuit die. We use machine learning methods to demonstrate FPGA operational state classification via magnetic field pattern correlation for both decapped and intact FPGA QDM images. This result provides an initial demonstration of functional integrated circuit characterization via magnetic field fingerprinting. Future work is required to determine whether and how this approach will be useful in areas such as integrated circuit security and failure analysis.

To date, the QDM's unique combination of magnetic field sensitivity, spatial resolution, field-of-view, and ease of use has allowed it to be used to measure microscopic current and magnetization distributions from a wide variety of sources in both the physical and life sciences [127–133]. Complementary to scanning techniques for characterizing integrated circuit magnetic field emanations, which include wire loops [134], probe antennas [135], magnetic force microscopy [117], SQUID magnetometers [113], and vapor cell magnetometers [136], the QDM employs a non-scanning imaging modality [121] that provides simultaneous high-resolution (micron-scale) and wide-field (millimeter-scale) vector magnetic imaging, while operating under ambient conditions. This capability allows for monitoring of transient behavior over sequential measurements of a magnetic field, providing a means to study correlations in signal patterns that can evolve more quickly than a single-sensor scan time. In addition, the QDM's simultaneous magnetic imaging modality is not subject to the reconstruction errors and drift that can arise from a scanned probe. With these distinctive advantages, the QDM technique is a promising approach for non-destructive physical testing of integrated circuits.

4.2 Experimental design

4.2.1 QDM experimental setup

A schematic of the QDM is shown in Fig. 4.1(a). The magnetic field sensor consists of a 4 mm \times 4 mm \times 0.5 mm diamond substrate with a 13 μ m surface layer of NV centers. The diamond is placed directly on the integrated circuit with the NV layer in contact with the integrated circuit surface. The diamond is grown by Element Six Ltd. to have an isotopically pure NV layer consisting of [^{12}C] \sim 99.999%, [^{14}N] \sim 27 ppm, and [NV^-] \sim 2 ppm. Light from a 532 nm, CW laser (Lighthouse Photonics Sprout-H-10W) optically addresses the NV layer with beam power of about 500 mW uniformly distributed over the 4 mm \times 4 mm NV layer. A flat-top beam shaping element (Eksma Optics GTH-5-250-4-VIS) and a cylindrical lens (Thorlabs LJ1558RM-A) create a rectangular beam profile (6 mm \times 6 mm) incident on the top face of the diamond at a sufficiently shallow angle of incidence (4°) relative to the top diamond surface to illuminate the entire NV layer. NV fluorescence is collected with a low magnification objective (Olympus UPlanFL N 4x 0.13 NA) to interrogate a large field-of-view of about 3.7 mm \times 3.7 mm. The fluorescence is filtered with a 633 nm longpass filter (Semrock LP02-633RU-25) and imaged onto a CMOS camera (Basler acA1920-155um). Resulting CW ODMR data is transferred to a computer where it is processed and analyzed with custom software utilizing LabVIEW and MATLAB.

A pair of 5 cm diameter SmCo permanent magnets (Super Magnet Man) is placed on opposing sides of the diamond to apply a uniform bias magnetic field (bias field) of $\mathbf{B}_0 = (B_X, B_Y, B_Z) = (2.0, 1.6, 0.7)$ mT to separate the resonances of the different NV axes (see Fig. 4.1(d)). X, Y, Z are the laboratory frame Cartesian coordinates with the X-Y plane defined as the surface of the diamond in Fig. 4.1(b). \mathbf{B}_0 induces a $\pm g_e \mu_B \mathbf{B}_0 \cdot \mathbf{n}$ Zeeman splitting of the spin triplet NV $m = 1$ and $m = -1$ ground states along each of four tetrahedrally defined NV symmetry axes, \mathbf{n} , with Landé g-factor g_e , and Bohr magneton μ_B . The hyperfine interaction between the NV and the ^{14}N isotope nuclear spin ($I=1$) results in an additional triplet level splitting. The four symmetry axes of the NV, shown in Fig. 4.1(b), are leveraged for vector magnetic field imaging using \mathbf{B}_0 projection onto all four NV

axes [137]. The ground state energy level diagram of a single NV axis is depicted in Fig. 4.1(c), neglecting hyperfine structure.

A 6 mm diameter copper wire loop made from 320 μm diameter magnet wire delivers 1 W of GHz-frequency microwave fields (TPI-1001-B and amplified with a Mini-Circuits ZHL-16W-43S+ amplifier) to drive the NV electronic spin transitions, $m_s = 0 \leftrightarrow -1$ or $m_s = 0 \leftrightarrow +1$, denoted by $f_{-,i}$ and $f_{+,i}$, respectively, see Fig. 4.1(c). The microwave power chosen is sufficiently low to not effect the observed normal function of the integrated circuit. The microwave field is modulated on and off through the use of a solid-state switch (ZASWA-2-50DRA+) controlled by a DAQ (NI-USB 6259) and synchronized with the frame acquisition of the camera to correct for laser intensity fluctuations and drift.

The intensity of optically induced NV fluorescence decreases for microwave fields on resonance with one of the spin transition energies. This decrease results from the $m = \pm 1$ spin-selectivity of the non-optical, intersystem crossing (ISC) mediated decay pathway for optically excited NVs [138]. The resonance frequencies between NV ground-state sublevels are determined from the ground-state Hamiltonian

$$H/h = (D(T) + M_z) S_z^2 + \gamma (B_x S_x + B_y S_y + B_z S_z) + M_x (S_y^2 - S_x^2) + M_y (S_x S_y + S_y S_x) \quad (4.1)$$

for the projection of \mathbf{B}_0 along a single NV axis, where h is Planck's constant, $D(T) \approx 2870$ MHz is the temperature dependent zero field splitting, T is the temperature, S_k are the dimensionless spin-1 Pauli operators, $\gamma = 2.803 \times 10^4$ MHz/T is the NV gyromagnetic ratio, B_k are the components of \mathbf{B}_0 in the NV frame, and M_k are crystal stress terms [139]. Electric field terms contribute minimally and are neglected [140–142]. Cartesian coordinates $k = x, y, z$ are defined in the NV frame with z along the selected NV axis, see Fig. 4.1(b). The contribution of the hyperfine interaction between the NV and ^{14}N nuclear spin is treated as a constant, 2.158 MHz energy level splitting and is not shown explicitly in Eqn. 4.1. Sweeping the frequency of the applied microwave fields across the range of resonant frequencies and collecting the NV fluorescence results in an optically detected magnetic resonance (ODMR) spectrum. Fig. 4.1(d) depicts the resulting ODMR measurements for

a bias field alignment where each NV axis experiences a different projection of the bias magnetic field.

Continuous-wave (CW) ODMR is used to image static FPGA magnetic fields. CW ODMR leverages continuous application of the laser and microwave field. This approach yields wide field-of-view images with high spatial resolution and good magnetic field sensitivity, while minimally perturbing the sample under study [121, 122]. A diamond with sufficiently low M_z inhomogeneity across the field-of-view is used to minimize degradation of performance [139]. Further suppression of strain contributions is achieved with application of the static bias field, \mathbf{B}_0 . Thus, the M_x and M_y terms in Eqn. 4.1 are negligible [139, 143]. The ground state Hamiltonian along a single NV axis reduces to

$$H/h \approx (D(T) + M_z)S_z^2 + \gamma B_z S_z + \gamma B_x S_x + \gamma B_y S_y, \quad (4.2)$$

and is used to determine the CW ODMR resonance frequencies for each pixel in a QDM image, and thereby to determine the magnetic field image from the sample FPGA.

4.2.2 Integrated circuit preparation, control, and layout

The Xilinx 7-series Artix FPGA (XC7A100T-1CSG324C) shown in Fig. 4.2(a) and 4.2(b) is selected for this study due to its versatility, general availability, and affordability. This FPGA is a 15 mm \times 15 mm wirebonded chip, fabricated in the TSMC 28 nm technology node, that has a \sim 6.5 mm \times 10 mm silicon die with eight clock regions. Digilent Nexys A7 development boards are used to operate and configure the Artix-7 FPGA. Two chips were used: one intact Artix 7 and decapsulated (decapped) Artix 7 that was prepared using a Nisene JetEtch Pro CuProtect decapsulator.

The large current draw and controllable location and size of ring oscillators make them ideal functional units for this study [144, 145]. Patterns of ring oscillators are implemented using the Xilinx Vivado Design Suite[®] to create distinguishable current distributions on the FPGAs for measure-

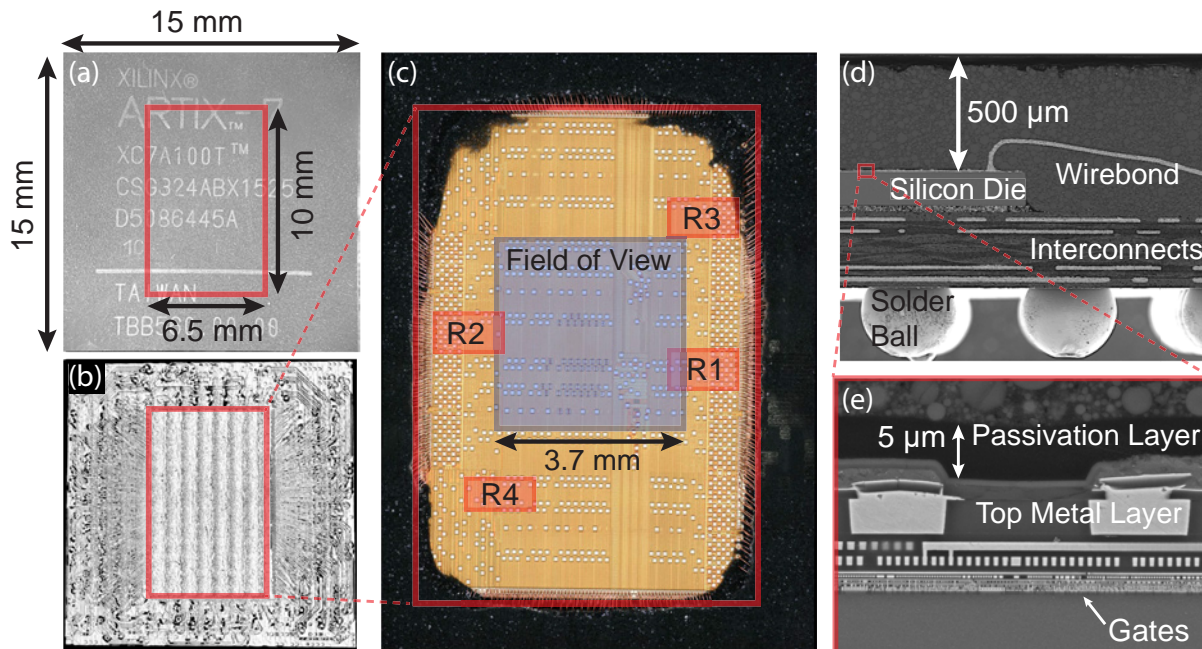


Figure 4.2: (a) Intact Xilinx 7-series Artix FPGA with die location and dimensions indicated in red (b) X-ray image of the FPGA package determining the position and size of die outlined in red (c) A high resolution image of the decapsulated FPGA with the fixed diamond measurement field-of-view indicated with a blue box, and the location of ring oscillator clusters indicated by red boxes labeled R1 - R4. (d) Scanning electron microscope (SEM) image of the FPGA cross-section showing the 500 μm stand-off distance between the chip die and the top layer. (e) Close-up of the SEM focusing on the metal layers of die. The thickness of the passivation layer is 5 μm and sets the minimum stand-off distance for the decapsulated measurements. (Figure credit: Matthew J. Turner)

ment by the QDM. Clusters of three-inverter ring oscillators are synthesized, placed, and routed to four different predefined clock regions on the FPGA, with clear spatial separation and spanning a majority of the die surface as shown in Fig. 4.2(c). A cross section of the die is shown in Figs. 4.2(d) and 4.2(e). The clusters consist of variable numbers of ring oscillators allowing for incremental increase or decrease of the current draw at the different locations on the FPGA. The active states of the FPGA are defined by ring oscillator clusters implemented in one of the predefined regions, and the idle state is defined as the FPGA powered on with no implemented ring oscillators. These active and idle states of the FPGA are used to create a lock-in type measurement of the chip activity. The ordering of states during a series of measurements are randomized to reduce susceptibility to systematic noise sources.

4.2.3 Experimental protocol and data analysis

CW ODMR measurements are taken with the FPGA in both active and idle states. The duration of each measurement is ~ 5 min per NV axis for each state. Such extended measurements are insensitive to transient effects on time scales shorter than the measurement; and sensitive to environmental drifts on the time-scale of the measurement. Magnetic field contributions from the ring oscillators are determined by subtracting the measured idle-state ODMR frequencies from the measured active-state ODMR frequencies, yielding the overall magnetic field due to the ring oscillators alone. For such measurements, the NV ground-state Hamiltonian is given by:

$$H/h \approx \left(D + \frac{\partial D}{\partial T} \Delta T + M_z \right) S_z^2 + \gamma(B_z + \Delta B_z) S_z + \gamma(B_x + \Delta B_x) S_x + \gamma(B_y + \Delta B_y) S_y \quad (4.3)$$

Here, terms with Δ originate from the FPGA active states and $\partial D/\partial T \approx -74$ kHz/°C [146]. Following these definitions and treating the off-axis magnetic fields as perturbative, the idle and active-state resonant frequencies for the upper (f_+) and lower (f_-) transitions of a single NV axis (i) are given by [143]

$$f_{\pm,i,Idle} \approx (D + M_z) + \frac{3\gamma^2}{2D} (B_x^2 + B_y^2) \pm \gamma B_z \quad (4.4)$$

and

$$f_{\pm,i,Active} \approx \left(D + \frac{\partial D}{\partial T} \Delta T + M_z \right) + \frac{3\gamma^2 [(B_x + \Delta B_x)^2 + (B_y + \Delta B_y)^2]}{2(D + \frac{\partial D}{\partial T} \Delta T + M_z)} \pm \gamma(B_z + \Delta B_z). \quad (4.5)$$

The desired FPGA state-dependent magnetic field projection on each NV axis, $\Delta B_{z,i}$, and the change in local temperature, ΔT , are given by

$$\begin{aligned} \Delta B_{z,i} &= \frac{1}{2\gamma} (\Delta f_{+,i} - \Delta f_{-,i}), \\ \Delta T &= \frac{1}{2\frac{\partial D}{\partial T}} (\Delta f_{+,i} + \Delta f_{-,i}), \end{aligned} \quad (4.6)$$

where $\Delta f_{\pm,i} = f_{\pm,i,Active} - f_{\pm,i,Idle}$. The off-axis magnetic fields of the sample are suppressed by the zero-field splitting; thus terms dependent on ΔB_x and ΔB_y are sufficiently small to be neglected in Eqn. 4.6. Terms dependent on B_x , B_y , B_z , D , and M_z are canceled by subtracting the idle resonance frequencies from the active state resonance frequencies. Determining the resonance frequencies from all four NV orientations for vector measurements, labeled by $i = 1, 2, 3, 4$ in Fig. 4.1(d), enables solving for the vector magnetic field in the lab frame

$$\begin{aligned} \Delta B_X &= \frac{\sqrt{3}}{2\sqrt{2}} (\Delta B_{z,2} + \Delta B_{z,4}), \\ \Delta B_Y &= \frac{\sqrt{3}}{2\sqrt{2}} (\Delta B_{z,1} + \Delta B_{z,3}), \\ \Delta B_Z &= \frac{\sqrt{3}}{4} [(\Delta B_{z,1} - \Delta B_{z,3}) - (\Delta B_{z,4} - \Delta B_{z,2})], \end{aligned} \quad (4.7)$$

where X, Y, Z are recalled to be the laboratory frame shown in Fig. 4.1(b).

The ODMR lineshape for NV ensembles is well described by a Lorentzian lineshape [143, 147].

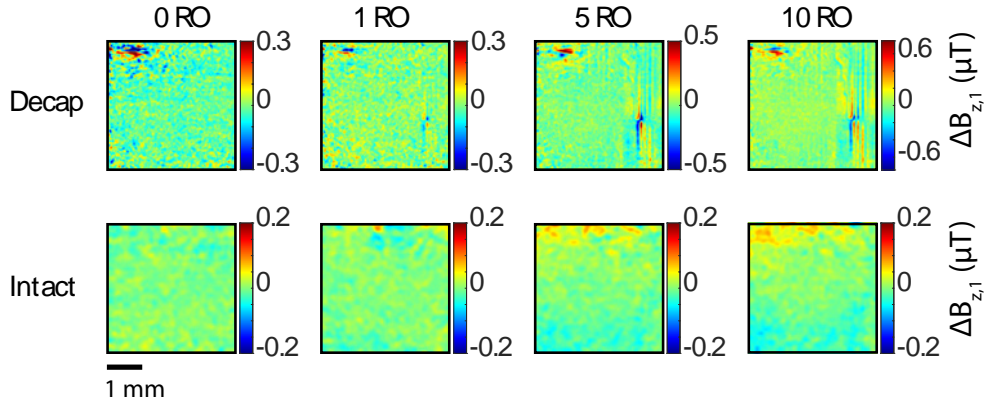


Figure 4.3: Representative QDM images of single runs of activity states with small numbers of RO that are used as inputs for the machine learning classification. The differences in the decapsulated chip states are observable through inspection, but the intact chip state differences are more subtle.

ODMR spectra for vector measurements of a ^{14}N diamond sample contain 24 resonance features, Fig. 4.1(d) (3 hyperfine features \times 2 electronic spin transitions \times 4 NV axes). The resonance frequencies of Eqn. 4.6 are extracted from the data by fitting all the Lorentzian parameters for every pixel in the field-of-view [122]. Furthermore, the contrast and linewidth of the resonances are determined, giving additional state dependent information which can additionally be used for probing high frequency magnetic fields [148]. GPU-based fitting algorithms [149] speed up this computationally intensive fitting and enable rapid analysis of a large number of measurements.

4.3 Data for low number active RO states

Representative QDM images taken from a series of measurements for single acquisitions of low numbers of ROs are given in Figure 4.3 for decapsulated and intact chips. Long term drift and large background fields are corrected for by subtracting off the nearest idle (0 RO) state in the measurement series. The measured states of the decapsulated chip are relatively easy to classify by visual inspection, consistent with the perfect accuracy in classification presented in Tbl. 4.1. The differences between states are more subtle in the intact data for low number of active ROs and will require machine learning to classify the different states as described in the following sections.

4.4 Machine learning analysis

Forward modeling of the current distributions and resulting magnetic fields for the different ring oscillator states programmed on the FPGA is an intractable problem without complete knowledge of the wire layout and current paths. However, automated machine learning algorithms can be applied to the QDM data to discriminate between and ultimately classify the different operating states. Ideally, a magnetic field image is used as input to a machine learning algorithm, and the functional state, defined as the number of active ring oscillators in this initial demonstration, is determined as the output. In practice, this problem is approached with a limited data set of magnetic field images for each FPGA state; and a dimensionality reduction algorithm is employed before applying a classification technique using the Python[76] package `scikit-learn`[150].

4.4.1 Data preprocessing

QDM data undergoes a series of preprocessing steps in preparation for dimensionality reduction and classification. Only images with Region 1 active are used so that the number of ring oscillators is predicted by the classification scheme. The number of ring oscillators activated for any given image is one of 0, 1, 5, 10, 50, 100, or 200. The data set consists of 40 QDM images per ring oscillator state for the decapsulated chip and 32 images per ring oscillator state for the intact chip. These $M \times N$ images are subsequently binned such that the decapsulated images contain 600×606 pixels and the intact images contain 300×303 pixels, while covering the same field-of-view. Measurements of the idle state (0 ring oscillators) are randomly taken during data collection to account for long term drifts. These idle state measurements are subtracted from active state images nearest in time. The intact and decapsulated data sets are split into training and test sets so that the prediction accuracy of the trained model can be estimated on data that the training procedure has not encountered. The splits are 75%/25% for the decapsulated images and 64%/36% for the intact images.

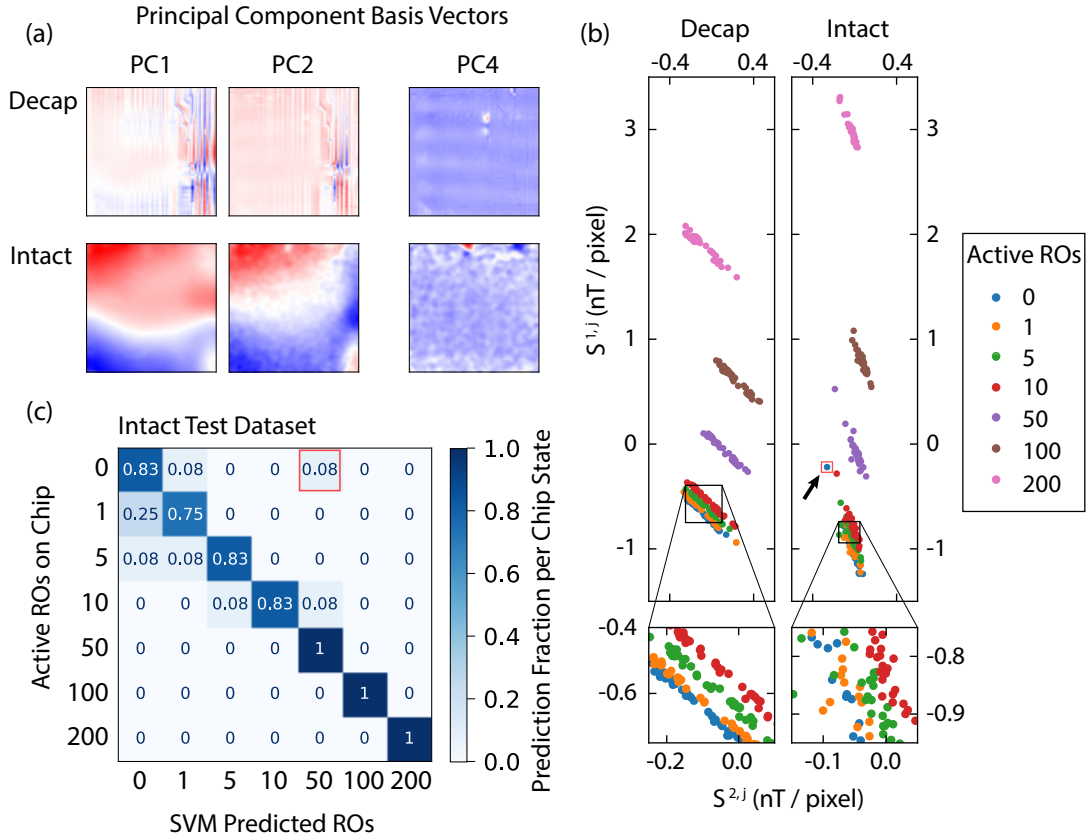


Figure 4.4: Principal component analysis (PCA) and support vector machine (SVM) classification of QDM images. Region 1 is active with 0, 1, 5, 10, 50, 100, or 200 ring oscillators. (a) Example principal component basis vectors plotted as images for both decapsulated (decap) and intact data sets. PC1 and PC2 are shown to exemplify principal components that resemble magnetic field images and thus will be useful in chip state classification. PC4 is shown as an example principal component that captures activity state independent variations and thus will not be useful in chip state classification. (b) The PCA score for PC1, $S^{1,j}$, is plotted against the score for PC2, $S^{2,j}$, for each magnetic field image, \mathbf{B}^j , as a demonstration of state distinguishability. This distinguishability is evidenced by the separation of colors representing differing numbers of active ring oscillators. Insets magnify the scores for small numbers of ring oscillators, and show greater fidelity of state separation in the decapsulated data set compared to the intact data set. (c) Table of SVM predictions on the test set for the intact images. For each unique value of active ring oscillators, there are 12 images in the test set. Rows indicate the fraction of images predicted for each of the possible chip states. All but one prediction (indicated by the red boxes in (b) and (c)) lie on or near the main diagonal demonstrating the high predictive power of the SVM classifier. The corresponding table for the decapsulated data set is not shown, as the main diagonal would contain 1's and the off diagonals would contain 0's due to the perfect separability of each state (see Table 4.1).




Figure 4.5: (a) The fraction of variance explained by the n^{th} principal component (PC) in the decapsulated dataset (green dots) and the intact dataset (blue triangles). (b) The cumulative fraction of variance explained by the first n PCs as a function of n (colors and markers are same as above). For both datasets, $> 99\%$ of the variance is explained by the first 9 PCs.

4.4.2 QDM image dimensionality reduction

Each magnetic field image is composed of $\sim 10^5$ pixels and thus occupy a high dimensional space for classification. Principal component analysis [151, 152] (PCA) is therefore used to reduce the dimensionality of the classification problem. PCA is a well-established technique that determines the highest variability axes of a high-dimensional data set. PCA amounts to an eigenanalysis where the eigenvectors, called principal components, correspond to the axes of interest; and the eigenvalues relate to the amount of data variance along the respective principal components.

PCA is applied separately to the data sets of the decapsulated chip and the intact chip with the `scikit-learn` class `decomposition.PCA()` and yields principal components such as those plotted in Fig. 4.4(a). Spatial patterns evident in the principal components are also present in the magnetic field images of Fig. 4.3, confirming that these features are physically significant and important for distinguishing between different samples. There exist as many principal components as dimensions in the data set; however, only the first several principal components capture non-noise based information. Figure 4.5(a) shows the fraction of variance explained by the n^{th} PC in both the decap-

ulated and intact datasets for the first 40 PCs. The first PC explains nearly all the variance in the magnetic field images, after which the explained variance drops off rapidly, ultimately depending only on noise and other RO state-independent information. Figure 4.5(b) shows the cumulative variance fraction for the first n PCs as a function of n . For the intact dataset, the first PC alone accounts for $> 99\%$ of the variance in the magnetic field images. For the decapsulated dataset, the first $n = 9$ PCs account for $> 99\%$ of the variance. The first $n = 9$ PC scores for each image are therefore used to capture all of the non-noise based information.

The scores of these first 9 principal components are used to effectively reduce the dimensionality of the magnetic field images from $\sim 10^5$ pixels to 9 scores. The principal component scores, $S^{i,j}$, are determined by taking the dot product of the i^{th} principal component, defined as \mathbf{W}^i , with the j^{th} image, \mathbf{B}^j , and normalized by the total number of pixels. This gives

$$S^{i,j} = \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N W_{m,n}^i B_{m,n}^j \quad (4.8)$$

for the first 9 principal components. Fig. 4.4(b) shows examples of the PCA scores: the score for PC1 is plotted against the score for PC2 for each magnetic field image of both the decapsulated and intact data (additional score plots are shown Fig. 4.6). The plot is color coded by number of active ring oscillators showing that these two scores are useful in distinguishing the number of active ring oscillators on the FPGA for both decapsulated and intact measurements. Classification of the active number of ring oscillators is accomplished by using the first 9 PCA scores as input to a support vector machine (SVM) classifier algorithm. The spread of data points along a fixed slope for each state in Fig. 4.4(b) is consistent with small offsets between different image acquisitions.

To further elucidate the relative importance of the first 9 PCs, Fig. 4.7 shows all of the first 9 PCs used in the chip state classification for both the decapsulated and intact datasets. A relative classification importance (defined below in section 4.4.3) is also calculated for each PC. The first two PCs of both datasets most strongly resemble the magnetic field images. As such, these PCs can naively be assumed to be the most useful for state classification, which agrees with these PCs having the highest importance factor. Additional PCs that visually capture state-independent information

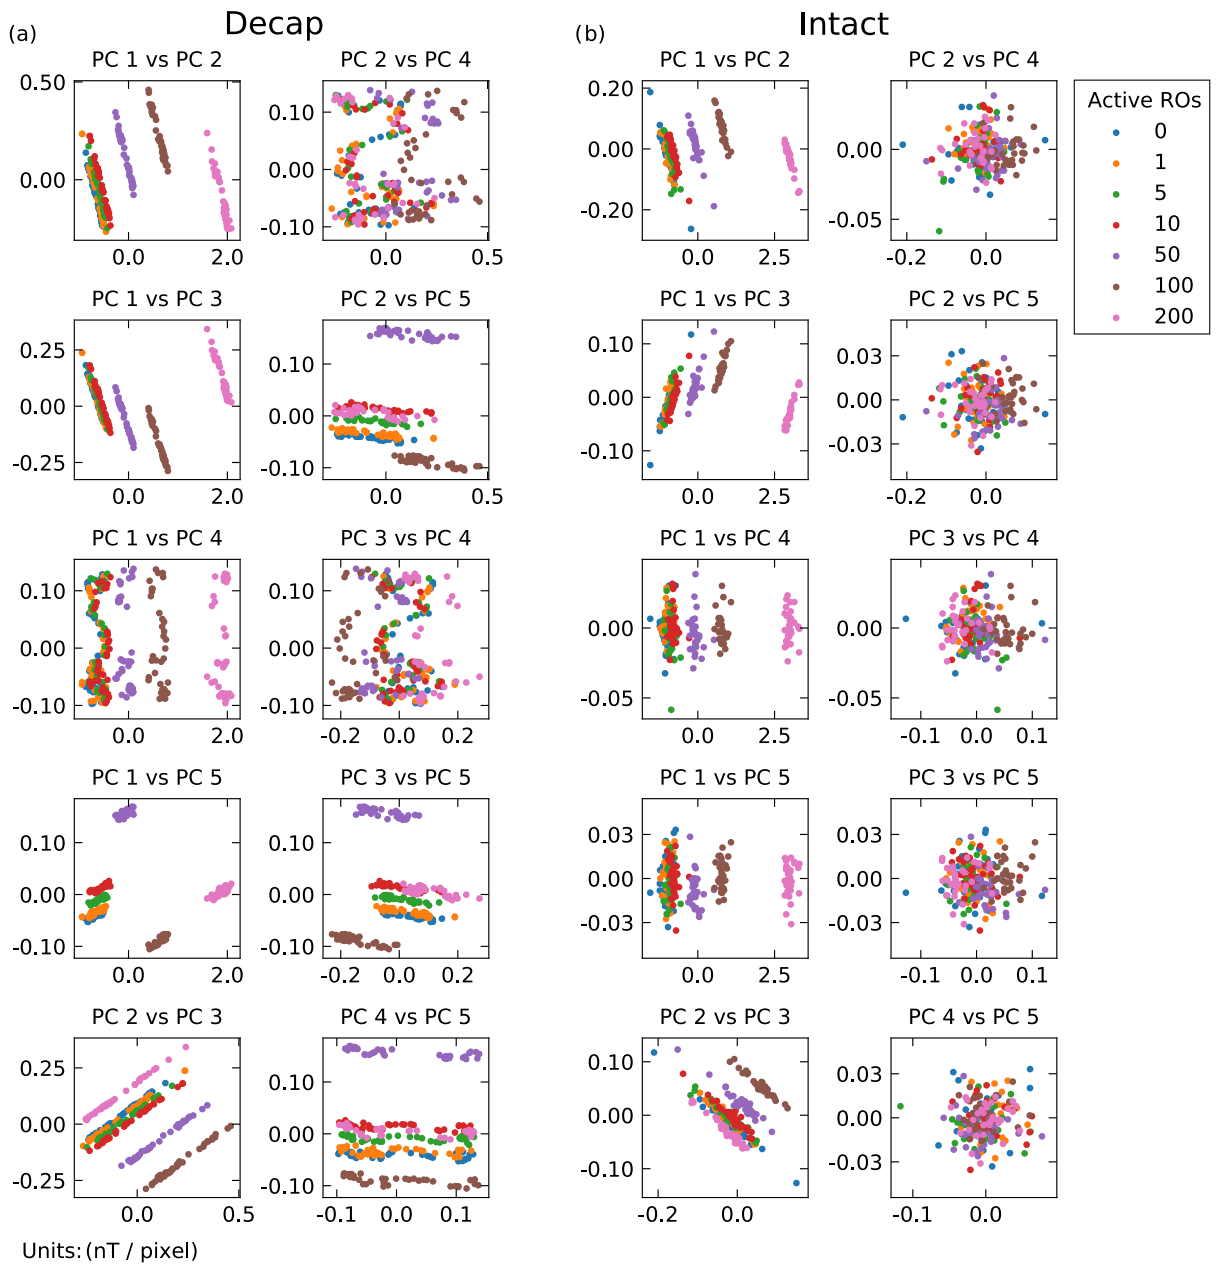


Figure 4.6: (a) Pairs of PC scores plotted against each other for PCs 1 through 5 in the decapsulated (decap) dataset. Units are nT / pixel and active number of ring oscillators (ROs) is color coded according to the legend. (b) The analogous plot for the intact dataset. Unlike Fig. 4.4(b), these plots are not on an equal aspect ratio and thus the slopes of groups of points carry less meaning.

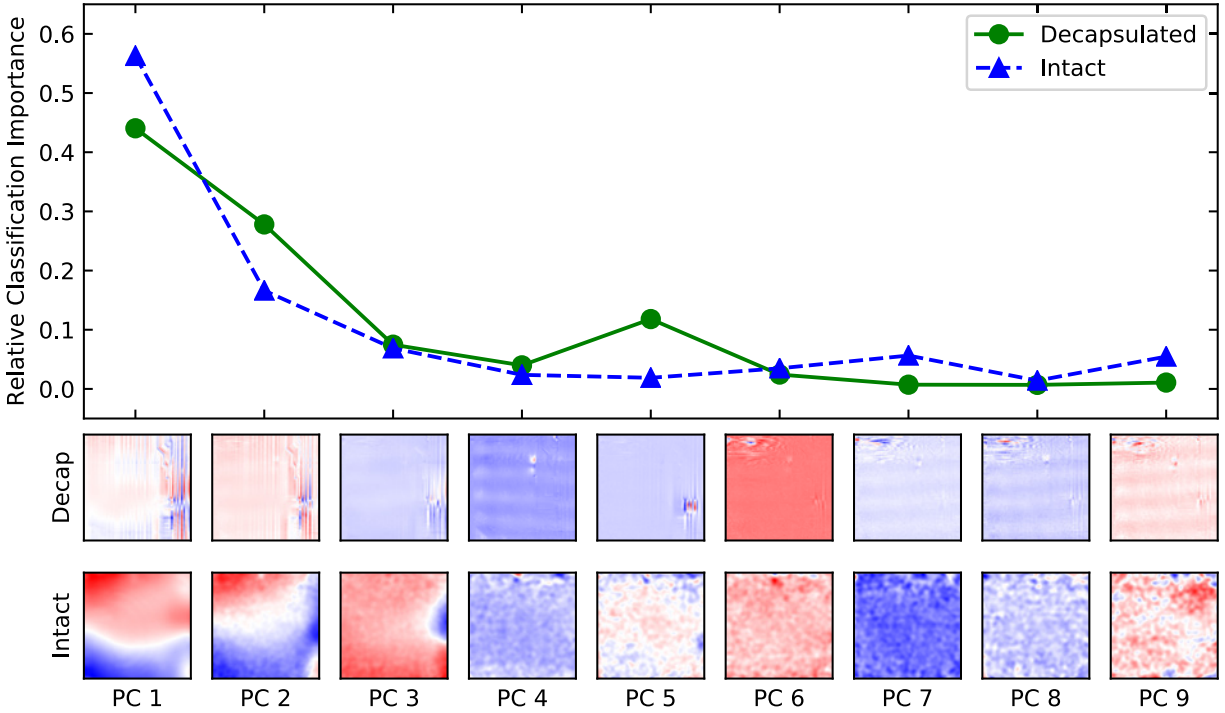


Figure 4.7: The relative importance for support vector machine (SVM) classification of each of the 9 principal components (PCs) used in both the decapsulated and intact datasets. The PCs that contain the most state dependent information have the most importance. Conversely, the PCs containing state independent information have the least importance.

accordingly have low importance. For example, PC 4 in the decapsulated dataset nicely captures the solder balls present in all magnetic field images, and PCs 6 through 9 capture activity in the upper left corner that is also present in all images.

In the intact dataset, PCs 4 through 9 appear noisy; however, PCs 7 and 9 both have a slightly elevated importance. The source of this importance is unknown, but given the small fields produced by the ROs on top of a larger background field, it is not surprising that PCs explaining a smaller fraction of data variance carry a higher importance in classification. The overall trend of decreasing importance with PC further strengthens the assertion that 9 PC scores suffice to classify each dataset.

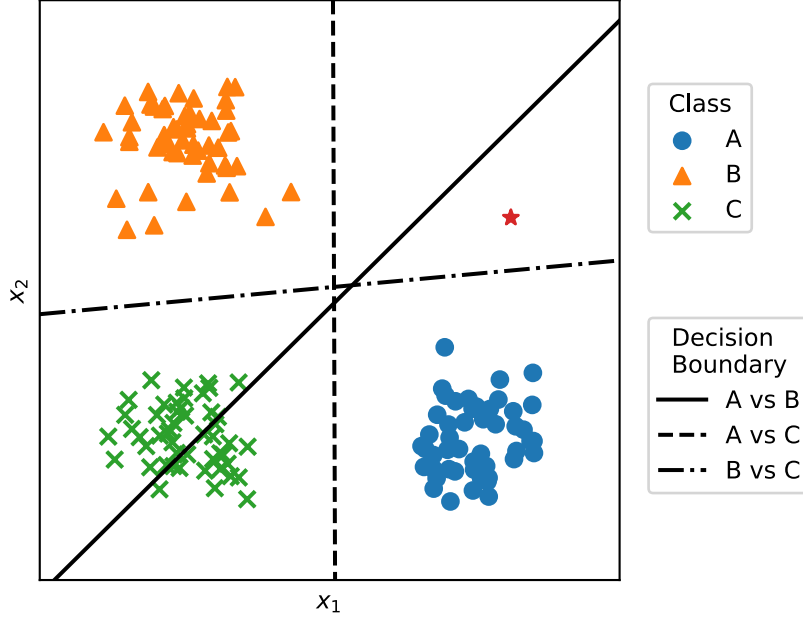


Figure 4.8: A simple SVM model (used only as a pedagogical example) with $N = 2$ dimensions (x_1 and x_2) and $M = 3$ classes (A, B, and C). Three lines represent the $K = 3$ hyperplane decision boundaries for each class comparison. The red star represents a new data point in need of classification. It will get two votes for A and one vote for B and thus will get a classification of A.

4.4.3 Integrated circuit activity state classification

A support vector machine [153] (SVM) is the supervised classification technique used to classify the magnetic field images, leveraging their key features characterized by the PCA scores. SVMs are a set of algorithms that seek to classify samples by creating a boundary between categories of a training data set that maximizes the gap separating those categories. Samples from a test set are then classified in relation to this boundary. A linear SVM seek to create a boundary between two classes, $y \in \{-1, 1\}$, of an N dimensional dataset by finding the $N - 1$ dimensional hyperplane that maximizes the orthogonal distance from both classes. Any hyperplane is defined by its normal vector and a constant offset. For the optimal hyperplane boundary defined by the SVM classifier, we denote its normal vector as \mathbf{w} and its constant offset as b . Any point, \mathbf{x} , lying in this hyperplane will satisfy the equation

$$\mathbf{w} \cdot \mathbf{x} - b = 0 \quad (4.9)$$

The SVM classification, y' , of an arbitrary point, \mathbf{x}' , will then be determined by whether this point lies above or below this plane, written mathematically as

$$y'(\mathbf{x}') = \text{sgn}(\mathbf{w} \cdot \mathbf{x}' - b) \quad (4.10)$$

We can interpret the normal vector, \mathbf{w} , as a weighting vector for each dimension in the N dimensional space and thus define a relative classification importance, I_j , of the j^{th} dimension as

$$I_j = \frac{|w_j|}{\sum_{j=1}^N |w_j|} \quad (4.11)$$

In the case of M classes (where $M > 2$), a one-versus-one classification scheme is implemented, whereby each class is compared with every other class. This yields $K = \frac{1}{2}M(M - 1)$ hyperplane boundaries, each with a normal vector, \mathbf{w}_k , and a constant, b_k , for $k = 1, 2, \dots, K$. A simple pedagogical example (not used for QDM data analysis) with $N = 2$ dimensions (x_1 and x_2) and $M = 3$ classes (A, B, and C) is shown in Figure 4.8. $N = 2$ gives lines for the hyperplanes, and $M = 3$ gives $K = 3$ decision boundary lines. The red star represents a hypothetical new data point in need of classification. It will get classified as A twice and B once. The red star will consequently receive a final classification of A.

The relative classification importance of the j^{th} dimension for the k^{th} one-versus-one comparison is defined as

$$I_j^k = \frac{|w_{k,j}|}{\sum_{j=1}^N |w_{k,j}|} \quad (4.12)$$

The overall relative classification importance of the j^{th} dimension will then be given by

$$I_j = \frac{\sum_{k=1}^K I_j^k}{\sum_{j=1}^N \sum_{k=1}^K I_j^k} \quad (4.13)$$

Equation 4.13 is used to calculate the relative importances of the 9 PCs shown in Fig. 4.7.

The `scikit-learn` class `svm.SVC()` is used as a multi-dimension, multi-category classifier for the QDM data. The categories for classification are the chip states given by the number of ring oscillators. The dimensionality is given by the 9 PCA scores recorded for each image. PCA scores are fit to the known FPGA states with a linear SVM model and a regularization parameter of $C = 6$.

When a training dataset contains outliers or is not linearly separable, the linear SVM classifier will struggle or even fail to find a hyperplane that will generalize well to predictions on the test set. To address this problem, implementations of SVM such as the `scikit-learn` class `svm.SVC()` used in this analysis will allow for some training samples to be misclassified in order to generalize well to the test set. The degree to which this is allowed is controlled by a regularization parameter, C . More precisely, a higher value of C attempts to correctly classify every training sample while a lower value of C will allow for more misclassifications of the training data. The regularization parameter thus represents a delicate balance between obtaining a good fit to the training set and generalizing well to the test set.

In order to determine the optimal value of C , we employ a 10-fold cross validation (CV) procedure. The training set is divided into 10 equal parts (folds). The first fold is used as a validation set and the remaining 9 folds are used as a new training set. A value of C is chosen and the PCA + SVM model is fit using the new training set. The model is then evaluated on the validation set and the accuracy is recorded. This process is repeated until each of the 10 folds have been used as the validation set. The CV accuracy for this value of C is then recorded as the average of the 10 validation accuracies. We repeat the entire process while varying C . Finally, we select the value of C that maximizes the CV accuracy and refit the model using all 10 folds as the training set. The test set containing as yet unseen data is then used to estimate the model accuracy by evaluating the final fitted model to this dataset.

The decapsulated dataset is easily separated by SVM and thus does not require CV of the regularization parameter. On the contrary, the intact dataset is not as easily separated and thus does require CV; figure 4.9 shows the results of this procedure. The discretization of the accuracies results from the relatively small training/validation sets (i.e. each jump in accuracy corresponds to

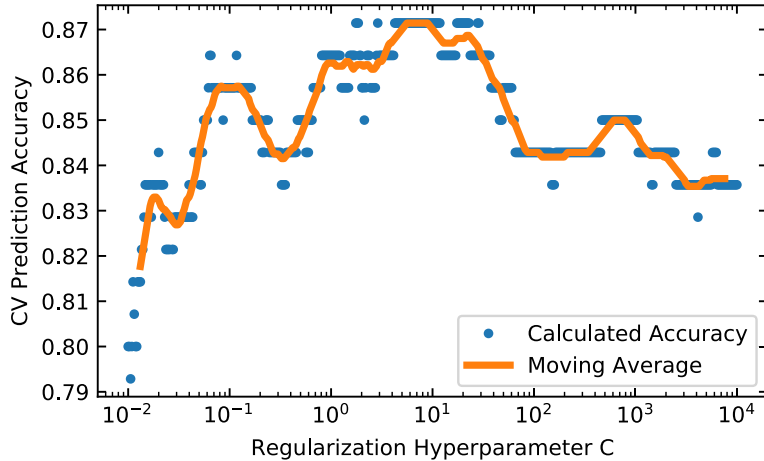


Figure 4.9: Cross validation of the SVM regularization hyperparameter, C . A 21 point moving average is included to aid the eye. The optimal value is selected at $C = 6$.

	Number of ROs (Region 1)							Total
	0	1	5	10	50	100	200	
Decapsulated	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Intact	0.83	0.75	0.83	0.83	1.00	1.00	1.00	0.89

Table 4.1: Chip state prediction accuracy on the test dataset (ROs = ring oscillators).

one more correct prediction in the validation set). A 21 point moving average is included to smooth out the CV prediction accuracy. It is evident from this curve that the optimal value of C occurs at about $C = 6$ and hence this value is chosen for the final model in the intact dataset.

4.4.4 Classification results

The full machine learning model, including preprocessing, PCA, and SVM, is fit using the training set and subsequently evaluated on the test set, for both decapped and intact FPGA data. The machine learning model efficacy, summarized in Table 4.1, is determined by the accuracy of the test set evaluations. FPGA activity states are well separated in PCA-space for the decapsulated data set. Predictions on the test set consequently yield perfect accuracy, even for small numbers of ring oscillators, consistent with expectations.

Results for intact data set are similarly well separated for large numbers of ring oscillators, resulting

in perfect prediction accuracy for ≥ 50 ring oscillators. However, FPGA activity states are not fully separated for < 50 ring oscillators, resulting in imperfect predictions. Nonetheless, the trained machine learning model achieves $\sim 80\%$ accuracy for each of 0, 1, 5, and 10 ring oscillator active states. Fig. 4.4(c) additionally shows that incorrect predictions are nearly always close to the correct state. For example, the model predicts 5 ring oscillators correctly in 83% of test cases, with misclassifications of 0 or 1 ring oscillator otherwise. The red box in Fig. 4.4(c) indicates a single case for which the classifier incorrectly predicts 50 ring oscillators for a 0 ring oscillator state. An arrow and analogous red box in Fig. 4.4(b) shows that the PCA score for this state is an outlier in the data.

The positive classification results presented in Fig. 4.4 give an initial demonstration of the capability of combined QDM and machine learning techniques to identify integrated circuit activity via non-invasive magnetic field imaging. The present results using ring oscillators are also potentially translatable to approaches that use the power side channel for chip activity analysis [154–156]. An extended dataset was taken to further elucidate the benefit of a large number of measurements per state on the prediction accuracy. Furthermore, in the extended dataset, the definition of state for the purpose of classification was expanded to include both number of ring oscillators and region activated. This dataset was also used to determine the impact of measurement stand-off distance on the ability to discriminate between spatially separated activity on the integrated circuit. The extended dataset was analyzed with the state classification criteria defined to include the spatial location of the active region in addition to the number of ring oscillators per region. For a simulated measurement stand-off distance of $\Delta z = 500 \mu\text{m}$, similar to that between the NV layer and FPGA current sources in the intact chip, a classification accuracy of $> 98\%$ is achieved on the extended dataset using 100 samples per state. This result indicates that a high degree of spatial state classification is possible in addition to local power classification using a large number of samples per state. Large data sets combined with more powerful machine learning methods have promise to enable classification of a wide array of chip activity in the context of hardware security and fault detection.

4.5 Extended dataset

An extended dataset of magnetic field images was collected in addition to the data presented above. This extended dataset further elucidates the impact of the number of measurements per state on the prediction accuracy, as well as the impact of measurement stand-off distance on the ability to discriminate between spatially separated activity on the integrated circuit. Two proximal regions, defined as Region 5 and Region 6, were used to further test spatial discrimination capabilities of PCA and SVM. The data consist of 100 measurements taken per state on the decapsulated chip for 12 different states. The states were composed of clusters of 0, 2, 4, or 6 ring oscillators that are activated in Region 5, or Region 6, or simultaneously in both Region 5 and Region 6.

The effect of measuring at different stand-off distances is simulated using upward continuation [120] to calculate the expected magnetic field at a large stand-off distance, Δz . Figure 4.10 shows the average of the 100 measurements for 3 of the 12 different states and projections at $\Delta z = 0 \mu\text{m}$ (decapsulated dataset), $50 \mu\text{m}$, $250 \mu\text{m}$, $500 \mu\text{m}$ (analog to an intact chip), and $1000 \mu\text{m}$.

Distinct spatial patterns are evident for the activation of 2 ring oscillators in region 5 and region 6. Specifically, a greater number of oscillations in the magnetic field polarity are present in region 6 compared to region 5. This difference is the result of different routing of power distribution in the top metal layer. Upward continuation of the data results in a decrease the magnetic field amplitude that occurs more rapidly in region 6 as a function of stand-off distance than in region 5. This difference is expected due to the cancellation of the large number of oscillating fields from neighboring wires in region 6. Even at $\Delta z = 1000 \mu\text{m}$, there is clear spatial desirability between all of the states, which is further quantified with PCA and SVM analysis.

The PCA + SVM analysis presented in the main text is performed on the extended dataset. However, in the case of the extended dataset, state classification is determined by the spatial location of the active region in addition to number of ROs per region. Thus a correct classification for the purposes of model accuracy consists of correctly identifying both the active region(s) as well as the

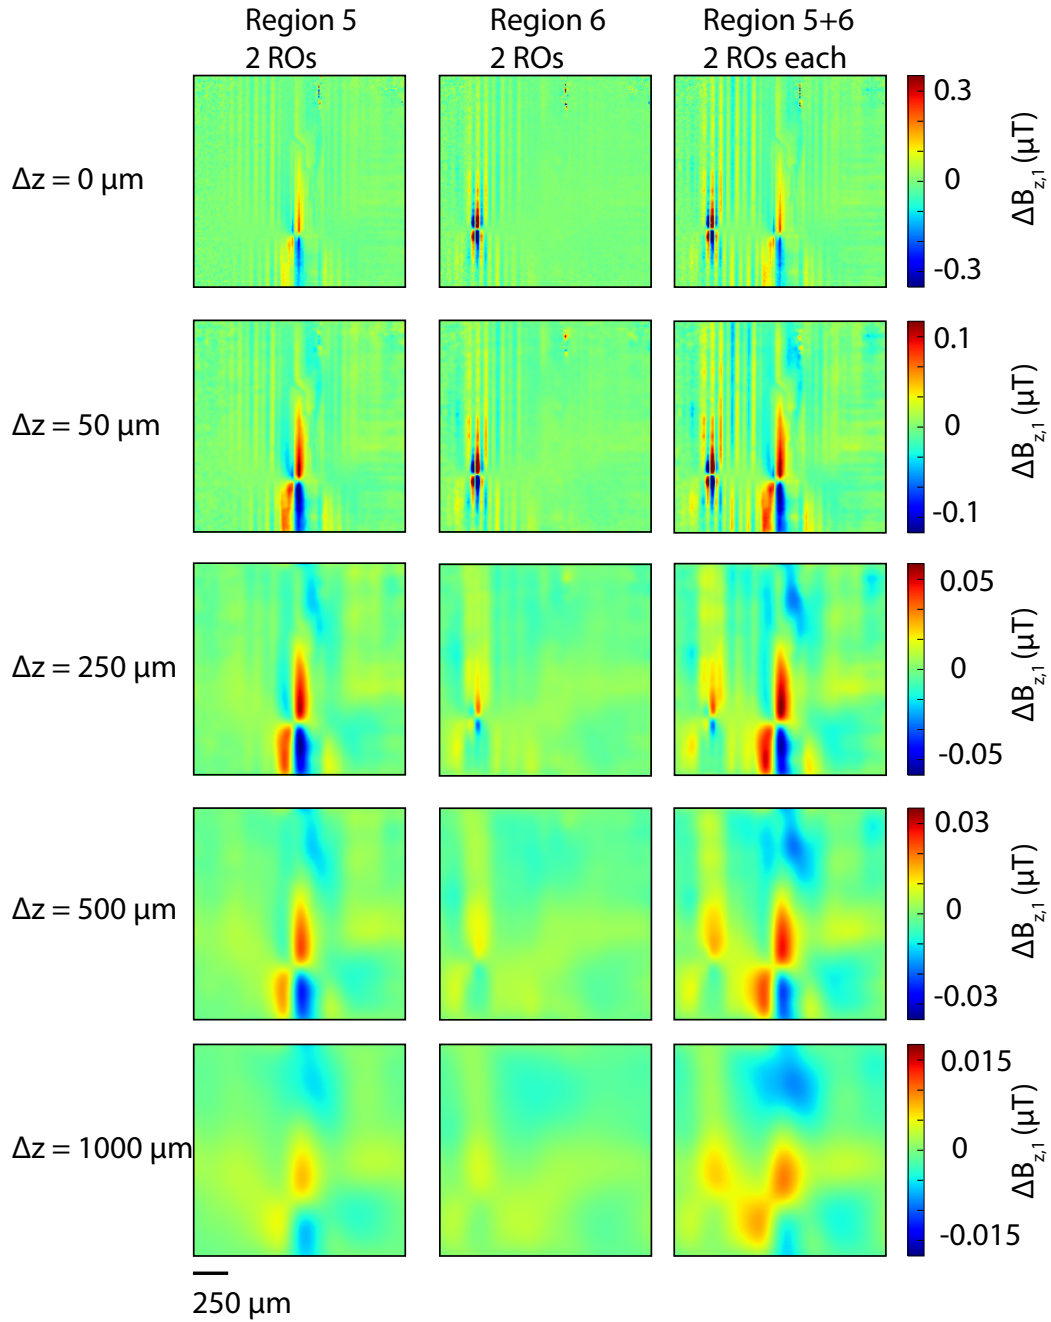


Figure 4.10: Subset of the extended QDM magnetic field dataset with 2 ring oscillators (ROs) active in region 5, region 6, or region 5 and region 6 simultaneously. The top row is high SNR data of 100 measurements taken from the decapsulated chip. Subsequent rows show the calculated magnetic image at different stand-off distances of $\Delta z = 50 \mu\text{m}$, $250 \mu\text{m}$, $500 \mu\text{m}$, and $1000 \mu\text{m}$. The $\Delta z = 500 \mu\text{m}$ row is the closest approximation of measurements taken of an intact chip.

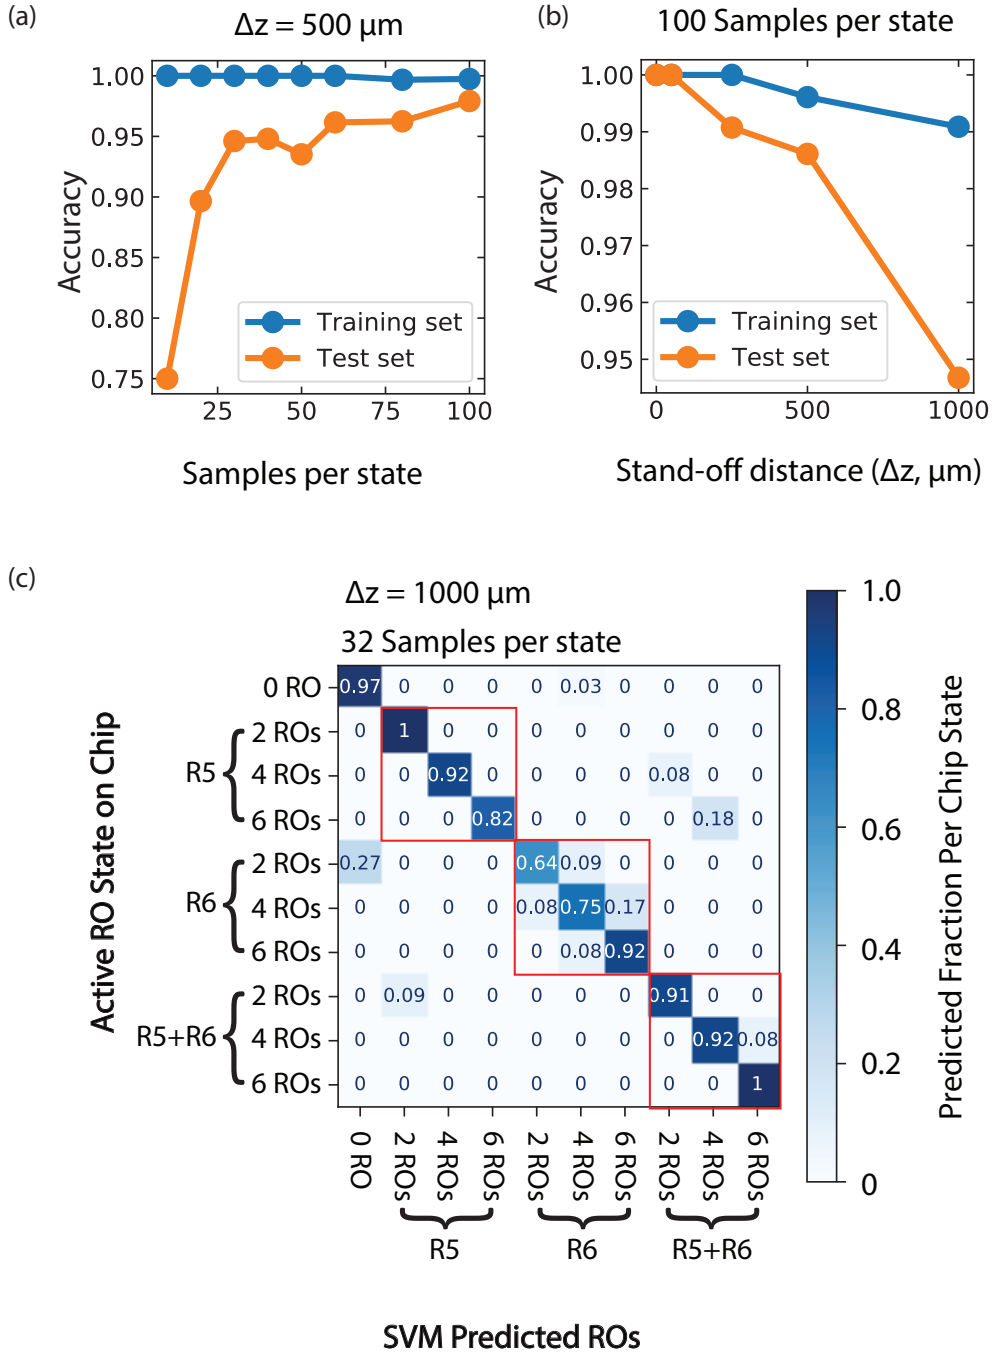


Figure 4.11: PCA + SVM model performance metrics of the extended dataset. (a) Model accuracy is plotted as a function of the size of the dataset for both the training and test sets. A stand-off distance of $500 \mu\text{m}$ is chosen to most closely replicate the original intact dataset. (b) Model accuracy is plotted as a function of the stand-off distance. The full dataset is used (100 samples per state). (c) Matrix of state predictions versus active state on chip at a stand-off distance of $1000 \mu\text{m}$. The matrix is row normalized to 1 so that each element represents the fraction of measurements of a given state that were predicted to be any state. The red boxes enclose predictions for which the predicted region and active region are the same.

number of active ring oscillators per active state; i.e. a state is defined as [active region, number of active ROs]. Accuracy is defined as the fraction of states correctly classified. The robustness of the dataset is tested by estimating the model accuracy as a function of the dataset size, shown in Figure 4.11(a). The upward continued dataset with $\Delta z = 500 \mu\text{m}$ is used and the size of the dataset is varied from 10 samples per state to 100 samples per state. A train-test split of 64%/36% is used to mimic the analysis of the original intact dataset. The training set has 100% accuracy when the number of samples per state is less than 60, but dips slightly below 100% for larger datasets. This is expected since the algorithm complexity is fixed but the absolute deviations in the data increase with sample size. Despite the increased error rate in the training set for larger samples per state, the larger datasets allow the model to learn the underlying structures of the dataset and consequently generalize better to unseen data. Thus the accuracy increases from 75% at 10 samples per state to 97% with 100 samples per state. Beyond 30 to 40 samples per state, the accuracy begins to level off, suggesting that 100 samples per state is enough data to get the maximum benefit from the PCA + SVM model.

The analysis is performed at various stand-off distances to test the predictive power of the model at varying levels of signal degradation. The difference between the intact chip and the decapsulated chip is about $500 \mu\text{m}$ of package material and thus $\Delta z = 500 \mu\text{m}$ is the best estimate for an analogous dataset taken from an intact chip. Figure 4.11(b) shows the model accuracy as a function of stand-off distance when the entire dataset is used (100 samples per state). As expected, perfect classification is achieved for small stand-off distances and monotonically decreases in accuracy with increasing stand-off distance. Importantly, for $\Delta z = 500 \mu\text{m}$, $> 98\%$ accuracy is achieved on the test set, implying that an analogous intact dataset would have a high degree of spatial state classification in addition to local power classification.

Finally, to further elucidate the ability to perform spatial classification, the state by state predication rates are shown in Figure 4.11(c). Results with 32 samples per state are shown to mimic the original dataset. A stand-off distance of $\Delta z = 1000 \mu\text{m}$ is chosen to show that prediction errors occur in accordance with expectations. Figure 4.11(c) presents the fraction of images with a given state that were predicted to be each possible state. Most predictions lie on the main diagonal as the overall

accuracy is about 89%. The red boxes show states from one region that were predicted to be in the same region. The majority of misclassifications are expected to lie in these red boxes. This is clearly the case, especially for region 6. In this region, 2 ring oscillators are misclassified as either 0 ring oscillators or 4 ring oscillators from the same region. 4 ring oscillators are misclassified as either 2 or 6 ring oscillators, and 6 ring oscillators are misclassified as 4 ring oscillators. The fact that this region has the highest error rate is consistent with expectation because the features in this region become less distinguishable as stand-off distance increases. Most of the remaining misclassifications are between region 5 and regions 5+6, which is expected since these regions look visually similar as seen in Figure 4.10.

Appendices

A

MCMC samples

This appendix contains plots of the MCMC samples drawn for the synthesized exoplanets shown in Fig. 2.4c. Each figure contains three trials for a 225 day period planet with a given baseline of data and semi-amplitude. We take 3 trials because sometimes the MCMC samples do not converge to the posterior distribution. Each row represents one of the five parameters fit with the Gaussian process (GP). The x-axis of each plot represents the MCMC sample number ranging from 1 to 1000. Each plot contains several colors, where each color represents the samples drawn from each walker in the MCMC, as described in [79].

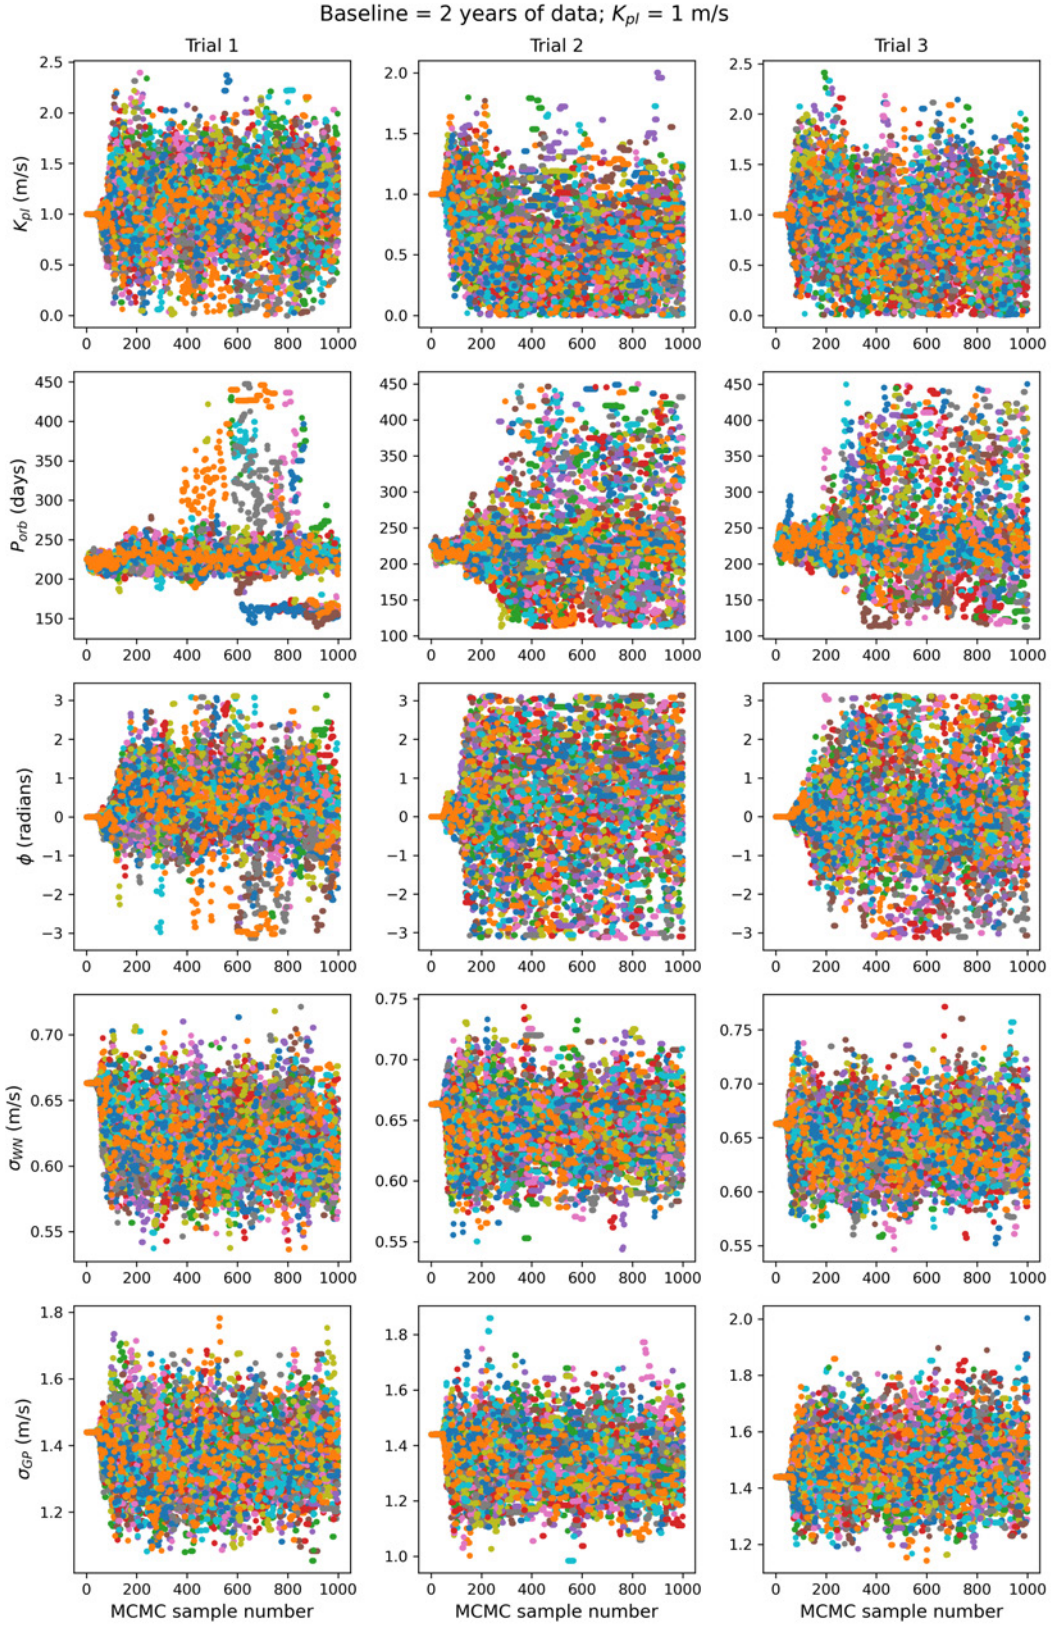


Figure A.1: MCMC samples for a synthetic 225 day planet with a semi-amplitude of 1 m/s. The data is synthesized with a baseline of 2 years.

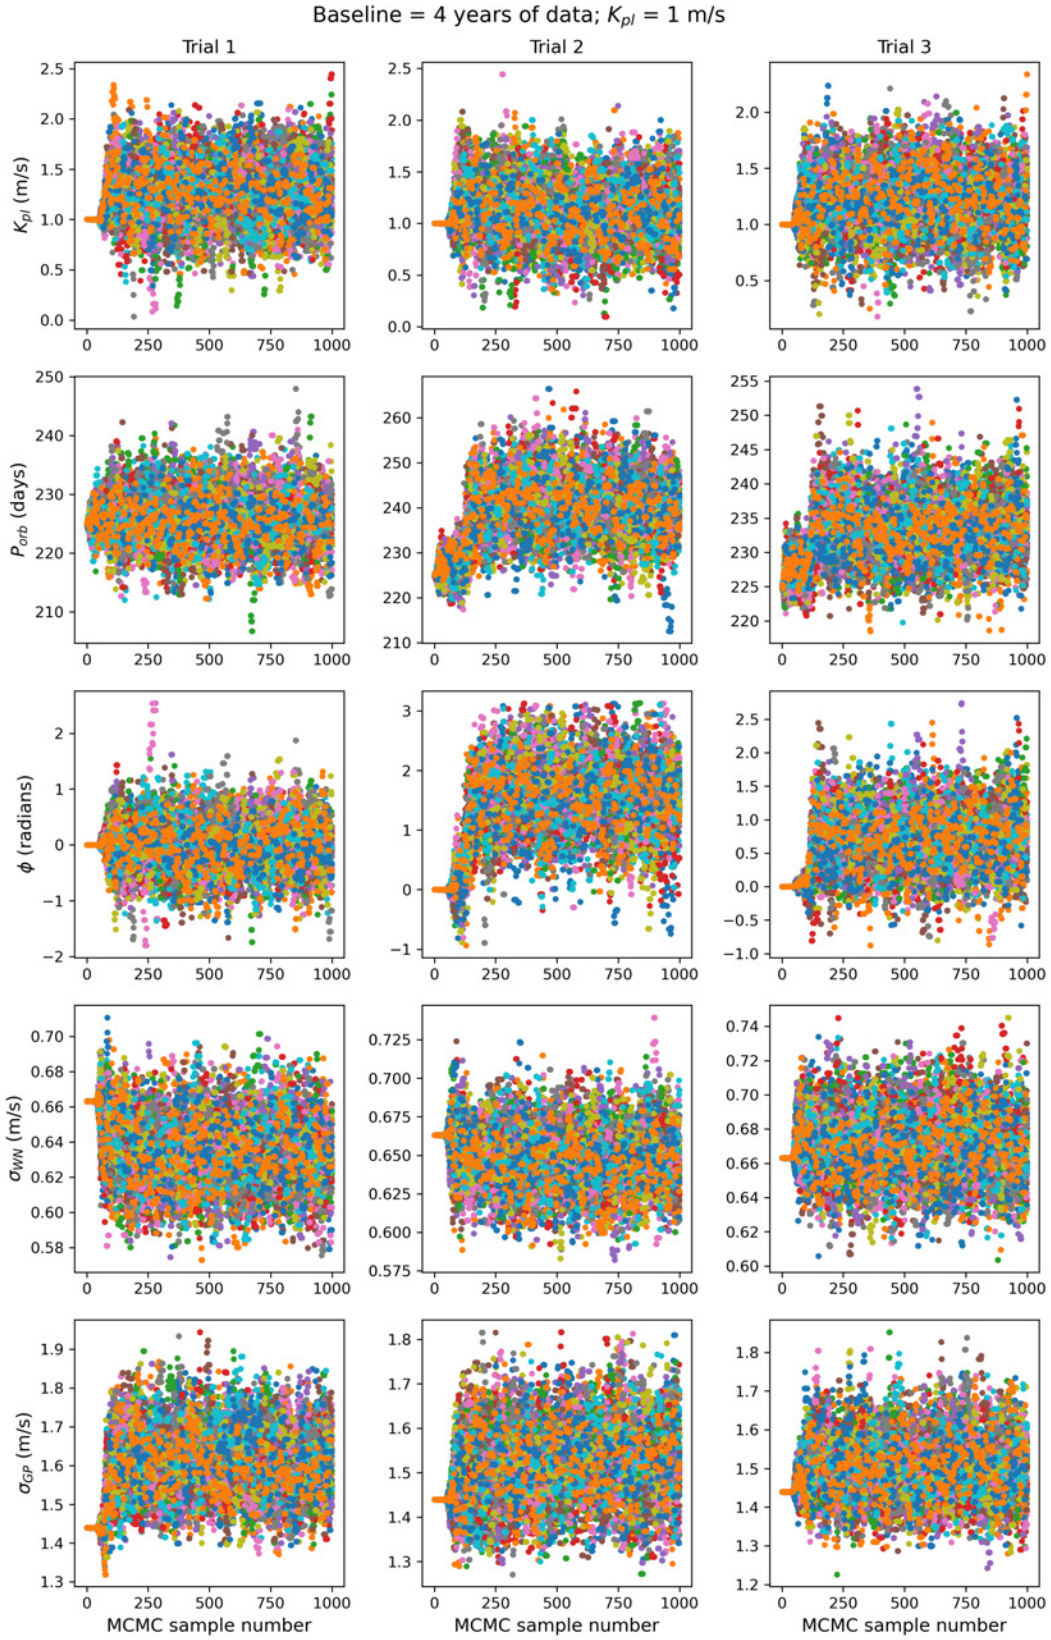


Figure A.2: MCMC samples for a synthetic 225 day planet with a semi-amplitude of 1 m/s. The data is synthesized with a baseline of 4 years.

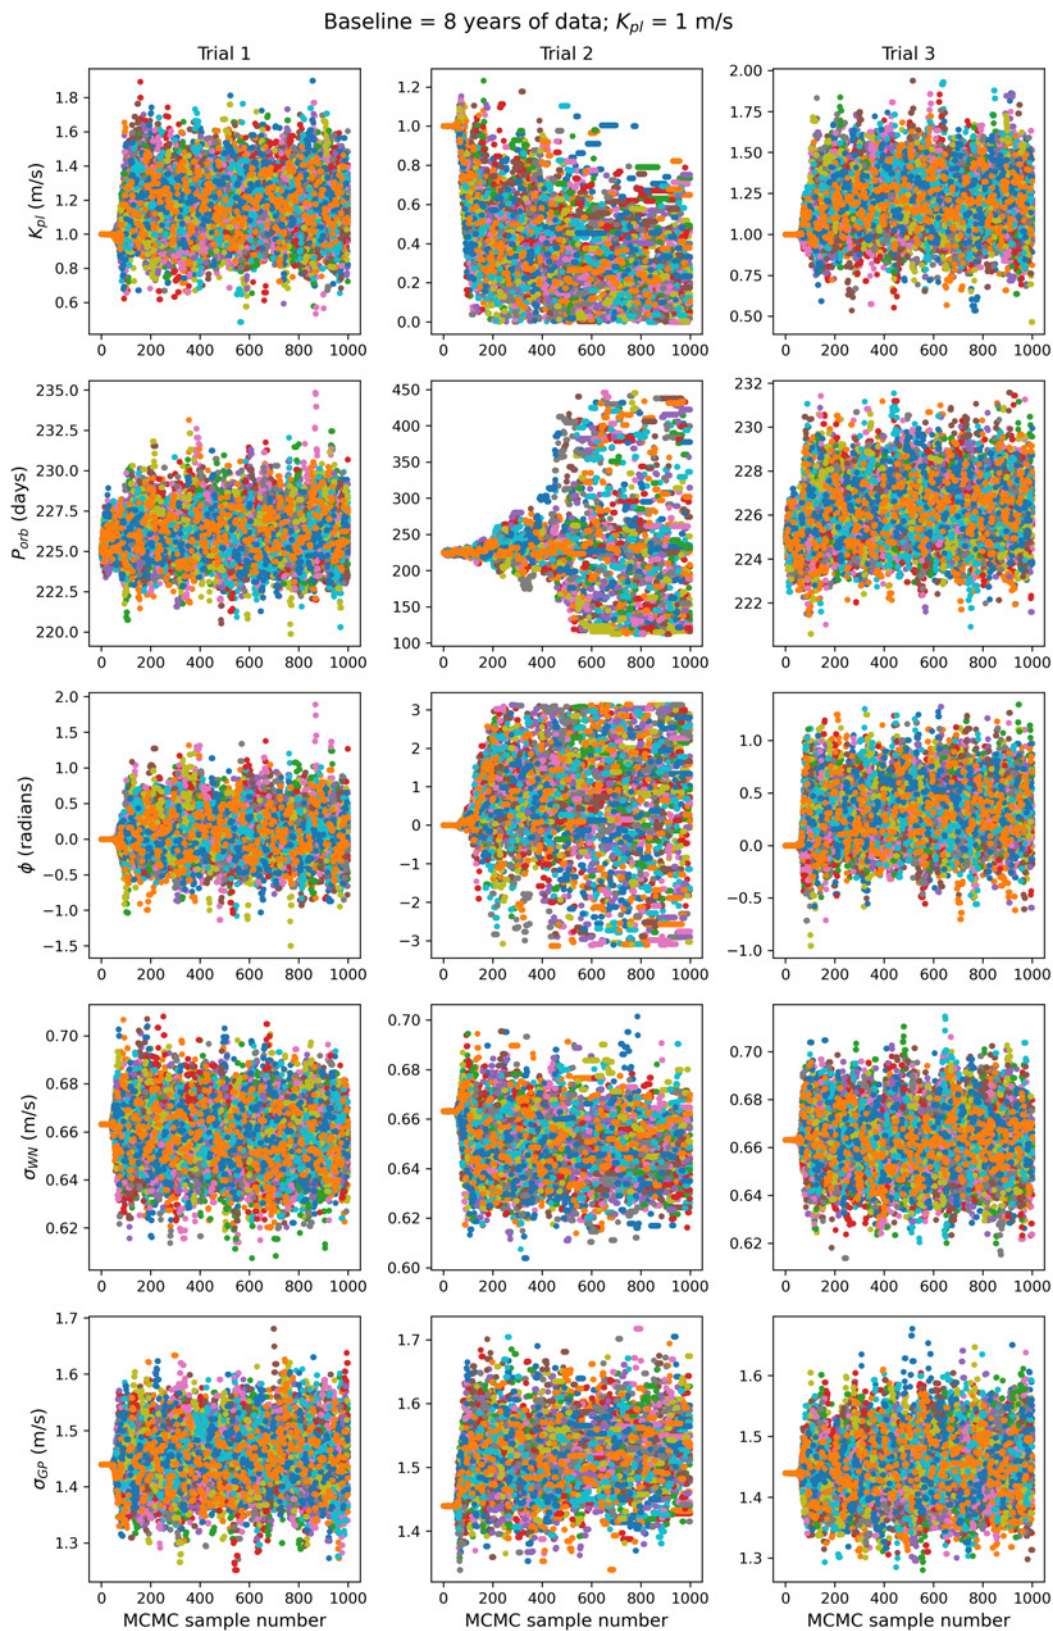


Figure A.3: MCMC samples for a synthetic 225 day planet with a semi-amplitude of 1 m/s. The data is synthesized with a baseline of 8 years.

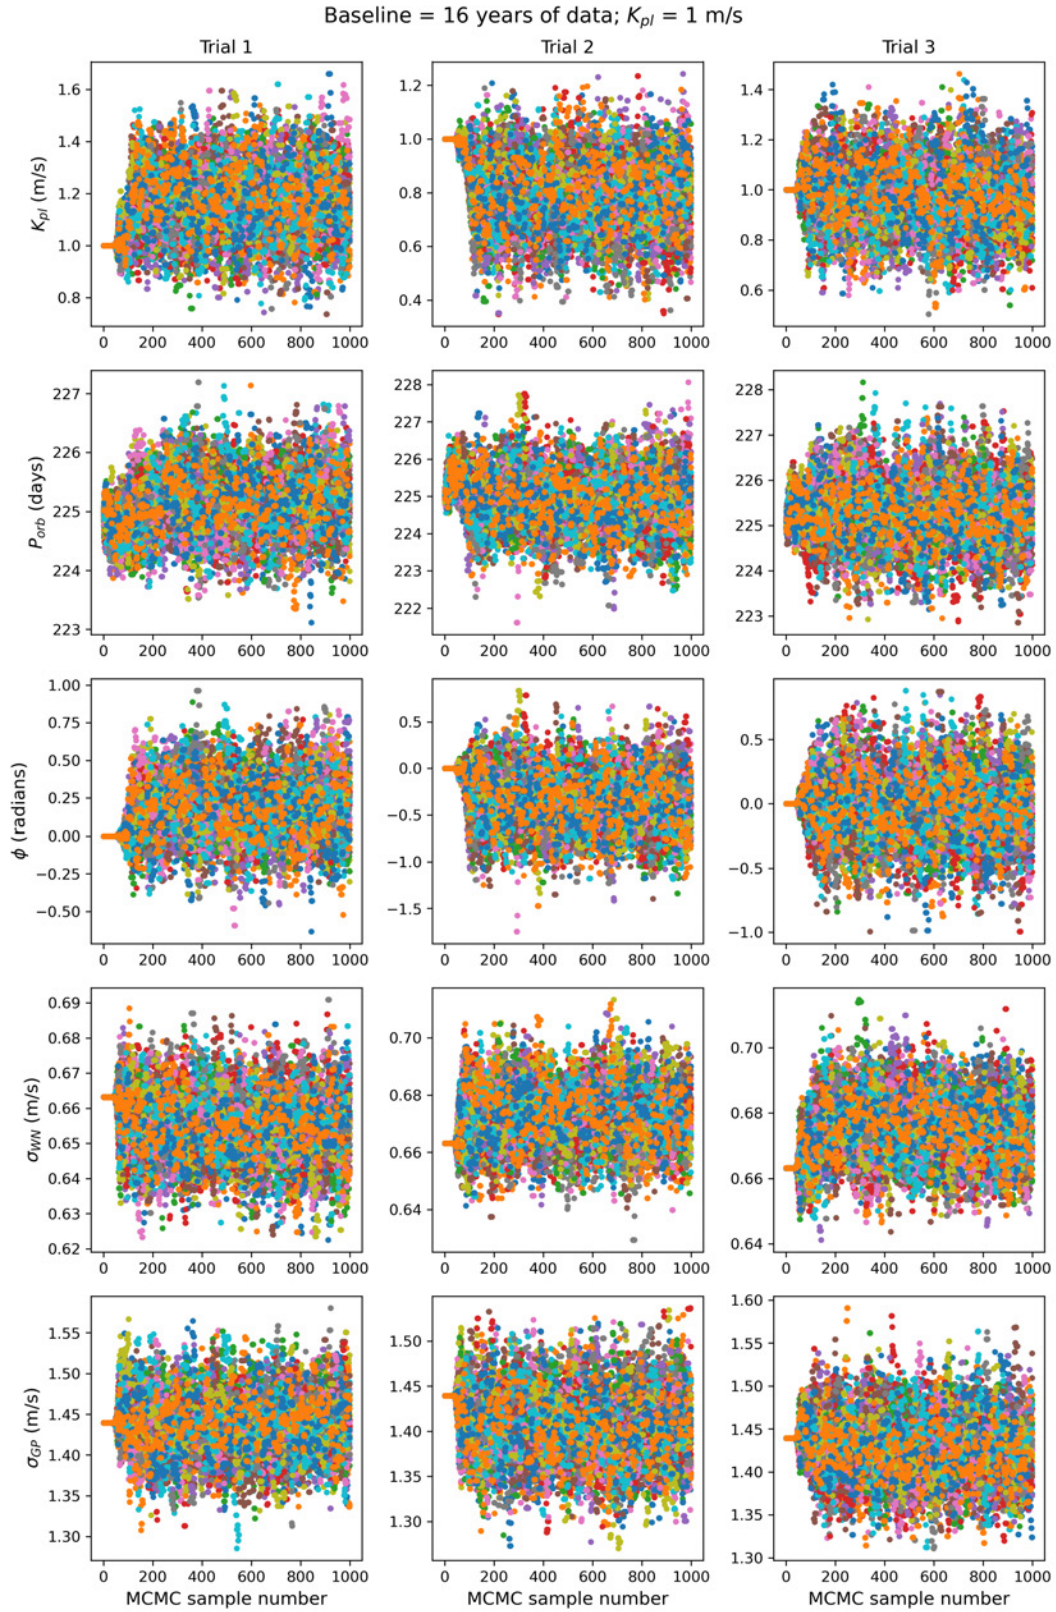


Figure A.4: MCMC samples for a synthetic 225 day planet with a semi-amplitude of 1 m/s. The data is synthesized with a baseline of 16 years.

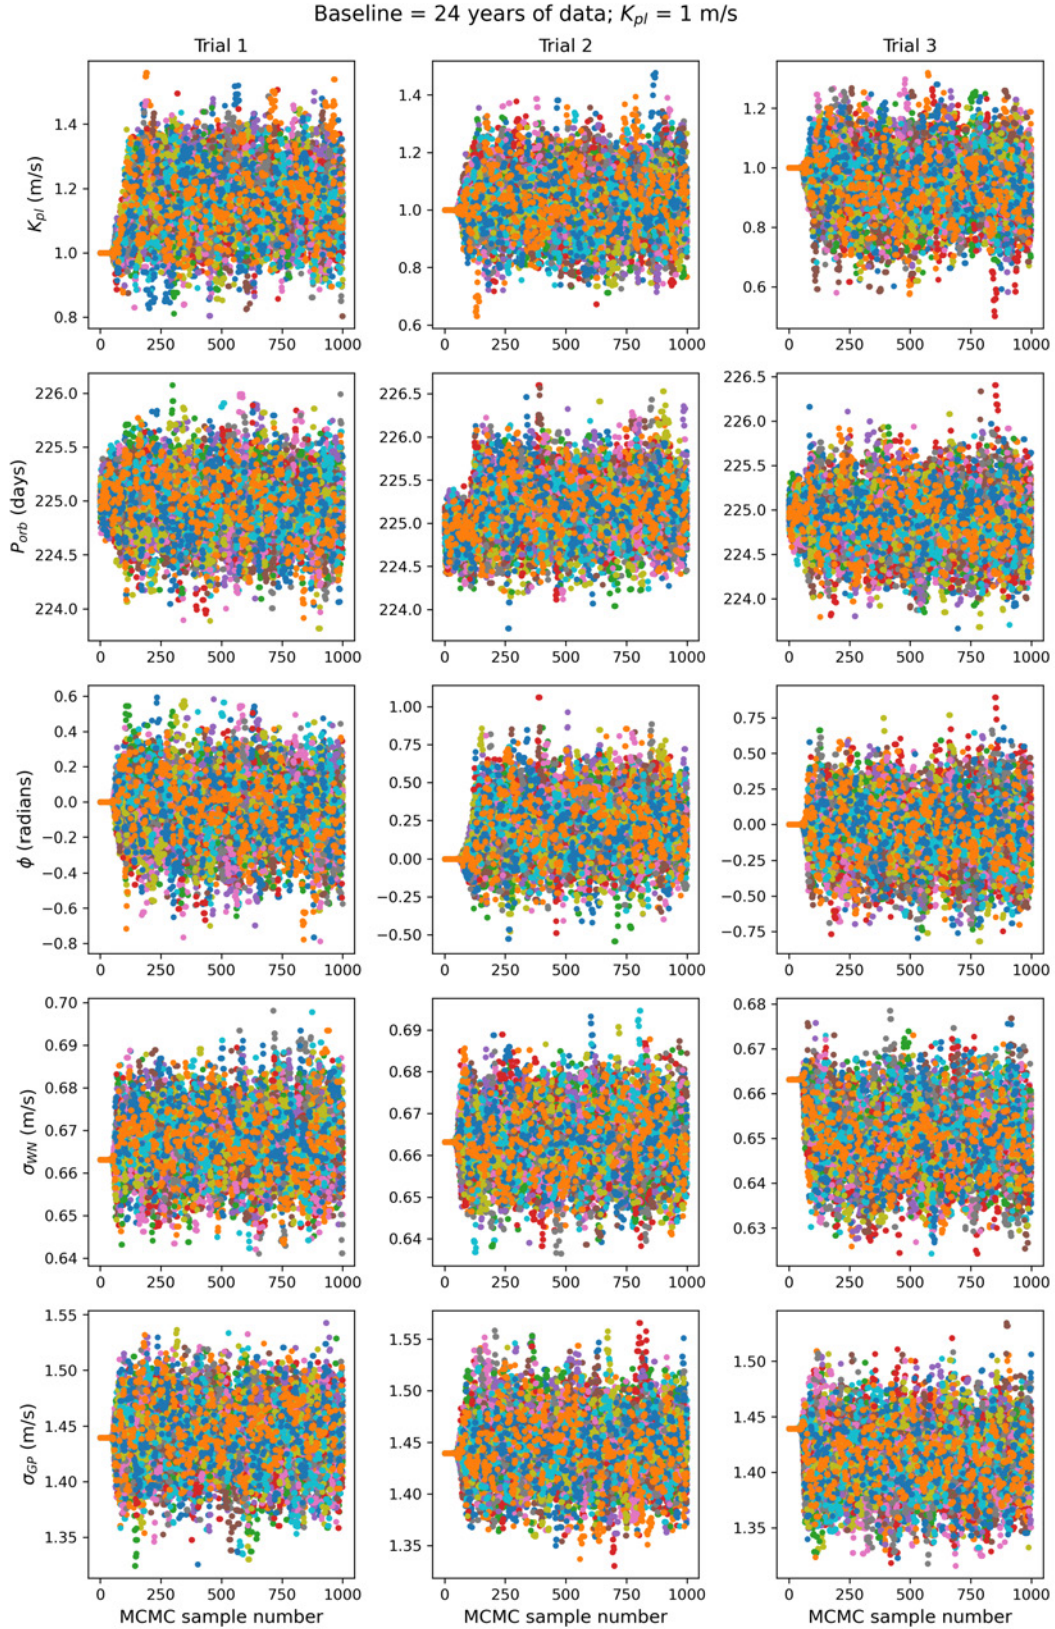


Figure A.5: MCMC samples for a synthetic 225 day planet with a semi-amplitude of 1 m/s. The data is synthesized with a baseline of 24 years.

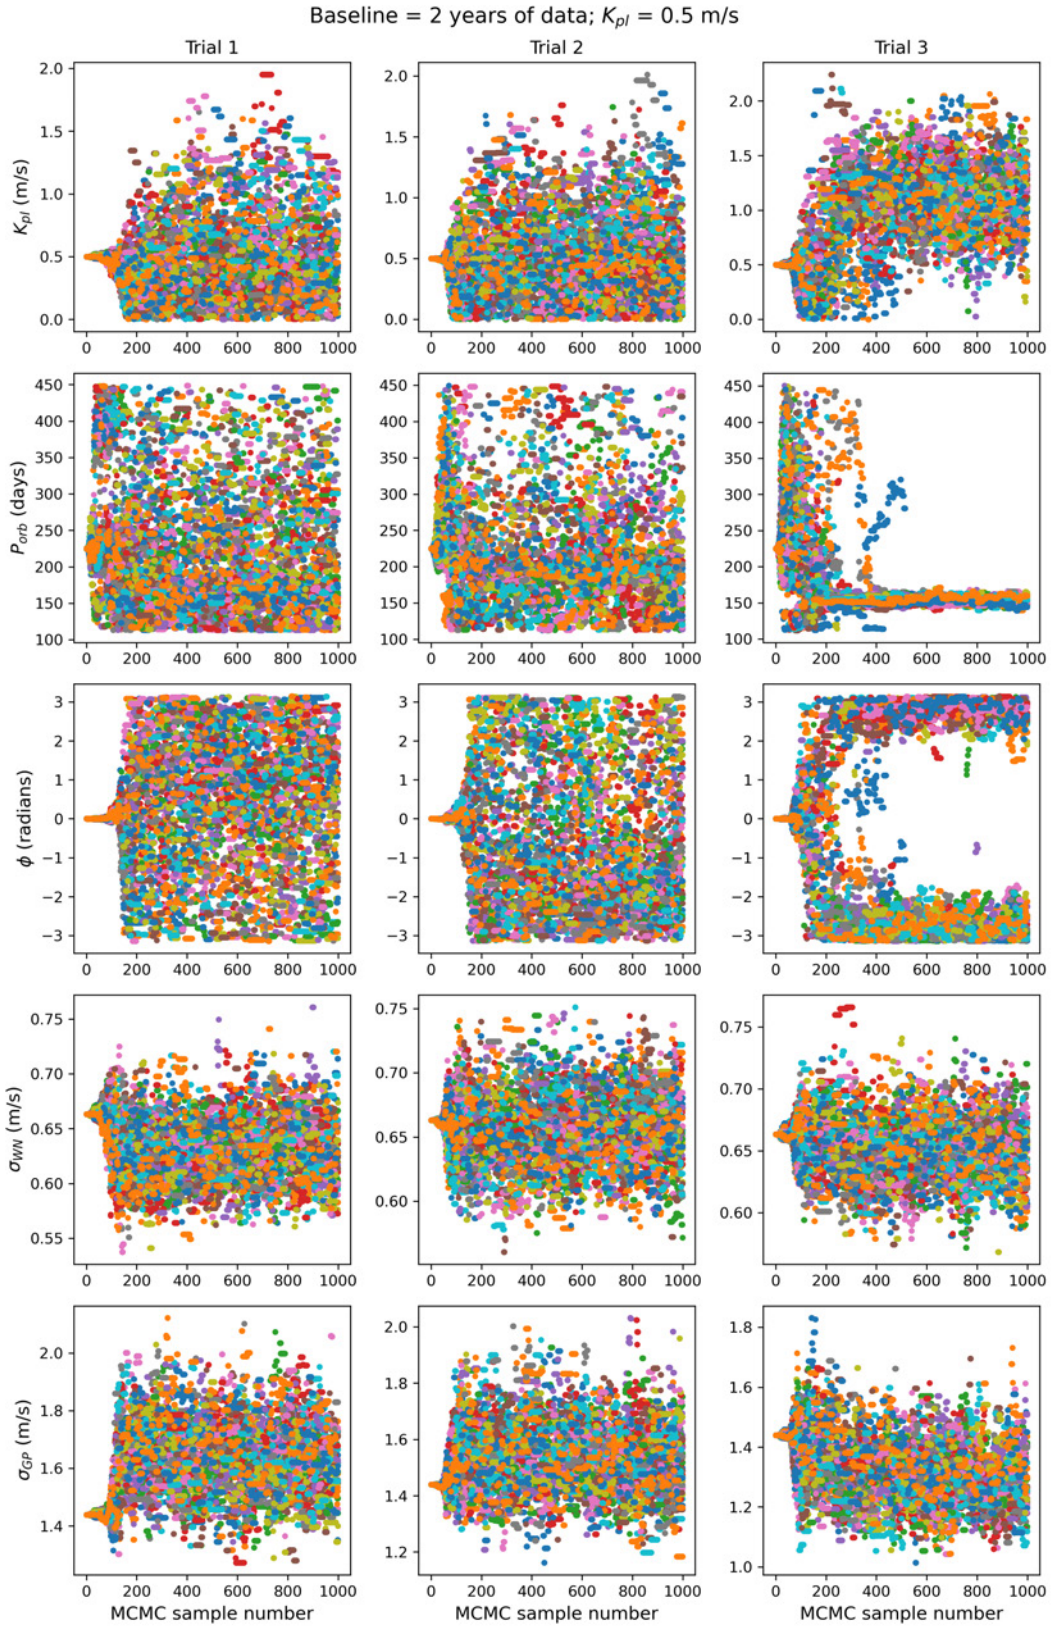


Figure A.6: MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.5 m/s. The data is synthesized with a baseline of 2 years.

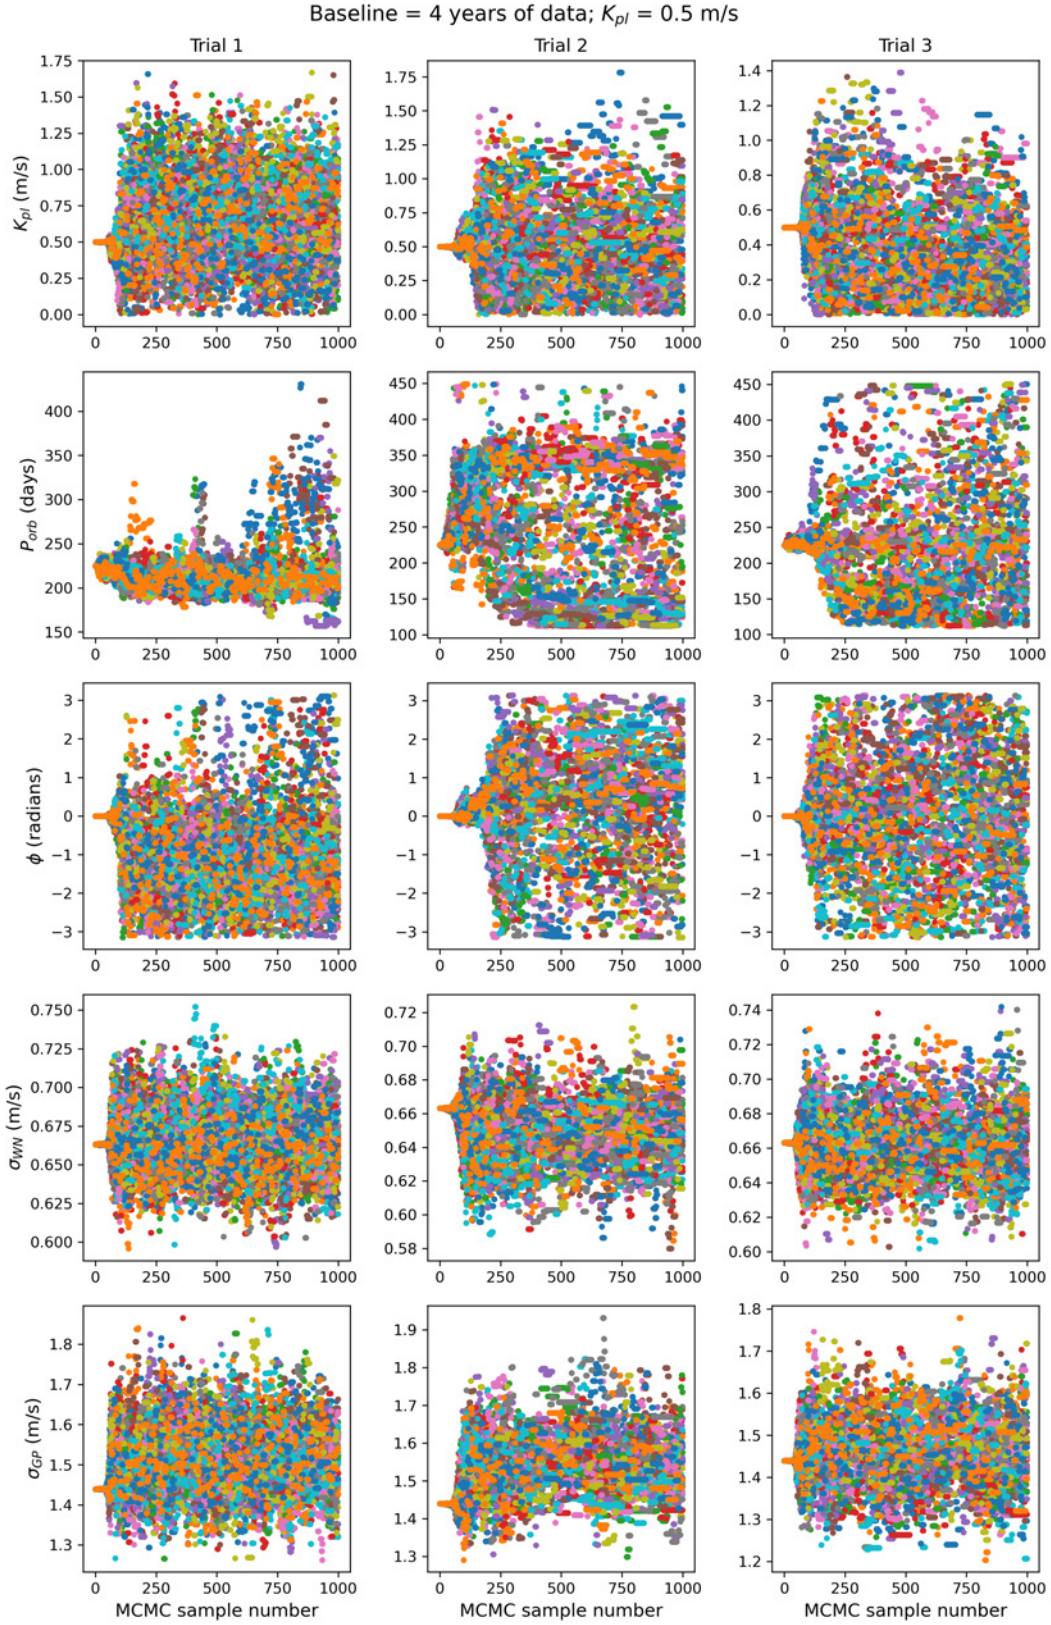


Figure A.7: MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.5 m/s. The data is synthesized with a baseline of 4 years.

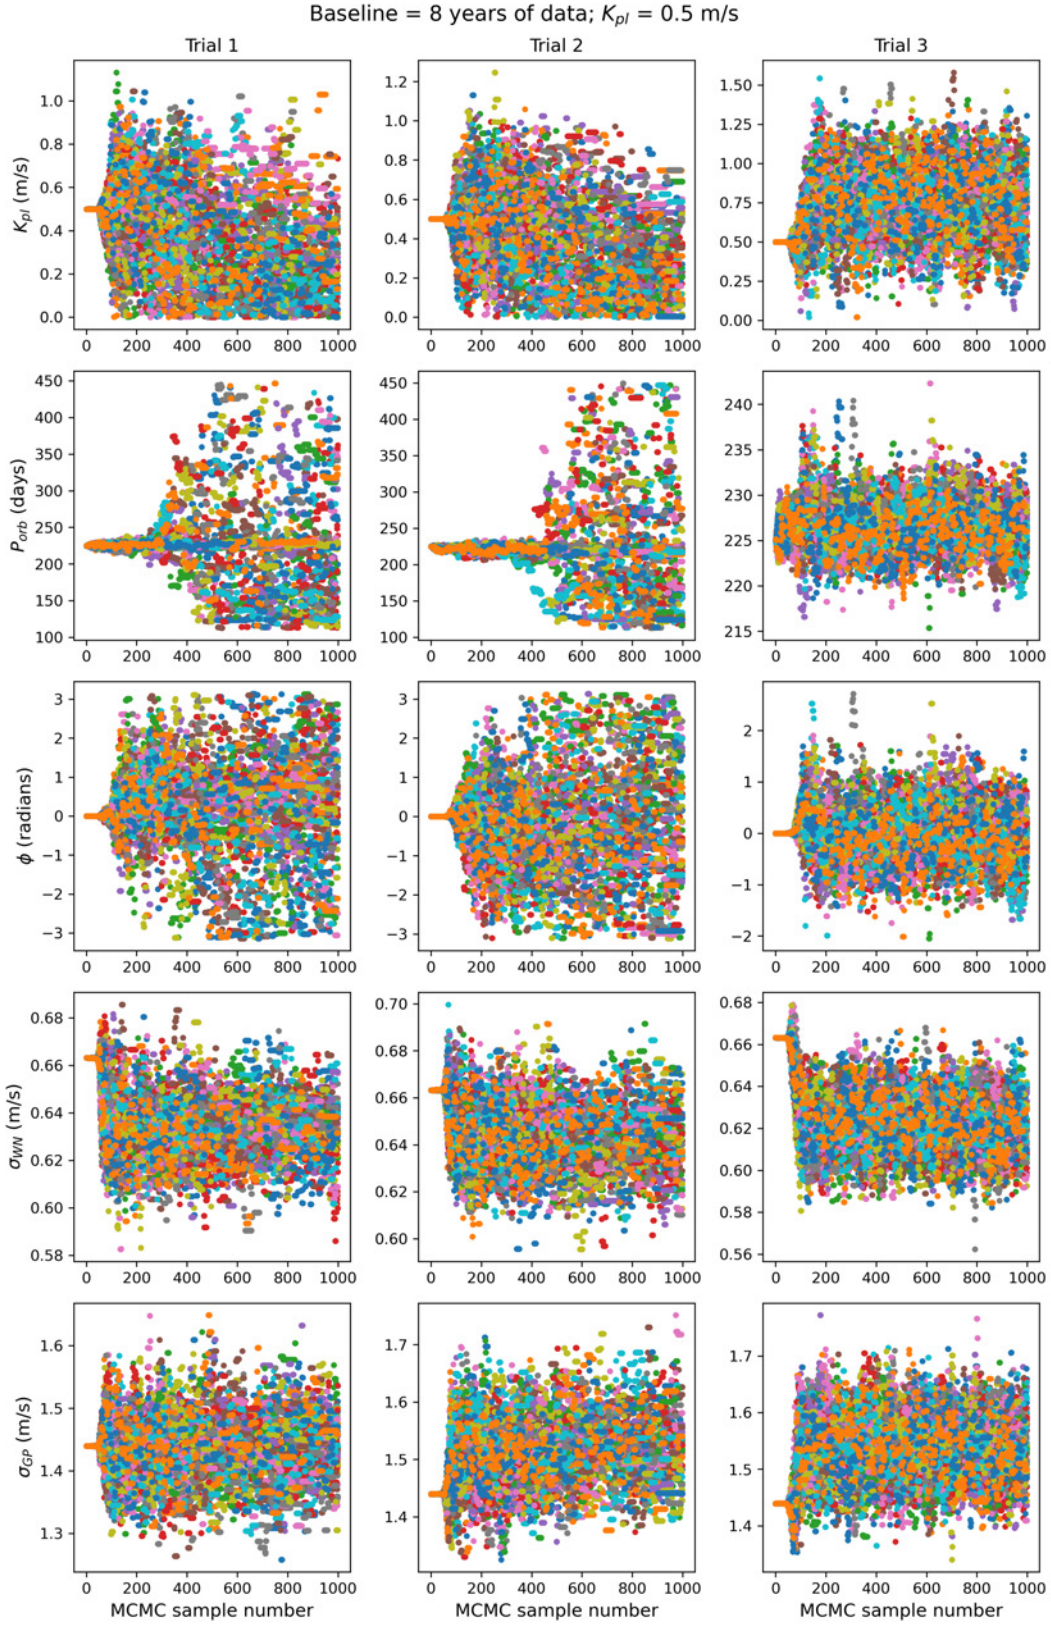


Figure A.8: MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.5 m/s. The data is synthesized with a baseline of 8 years.

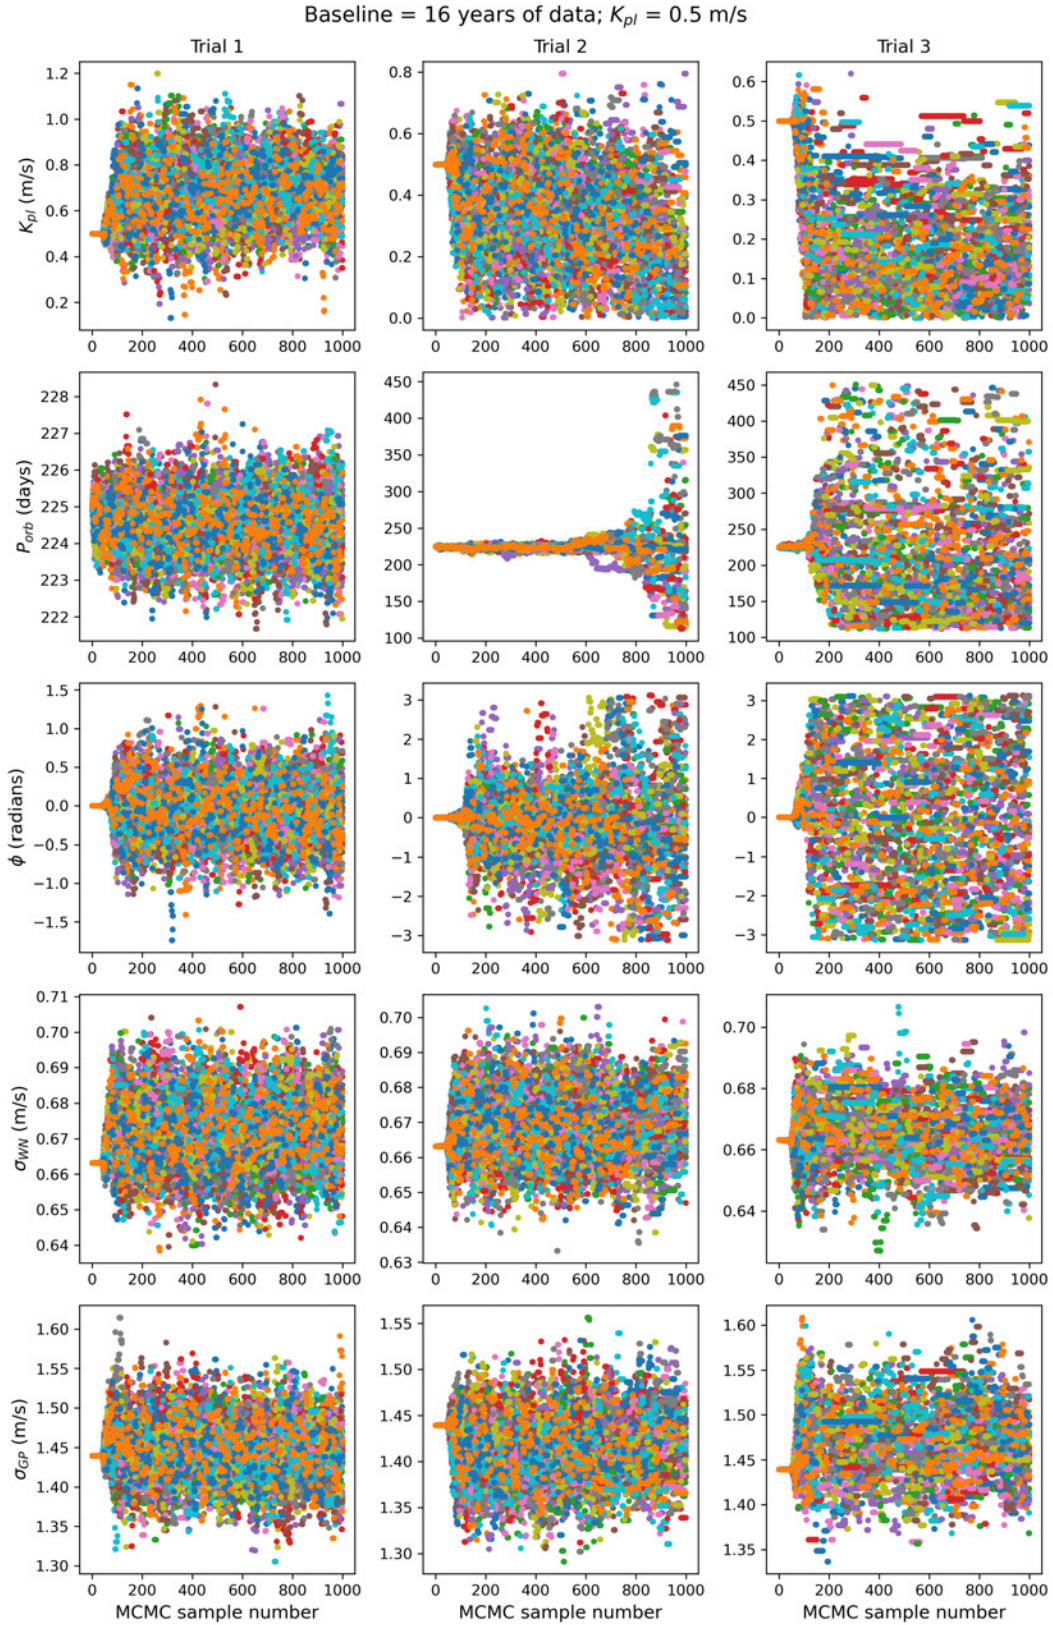


Figure A.9: MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.5 m/s. The data is synthesized with a baseline of 16 years.

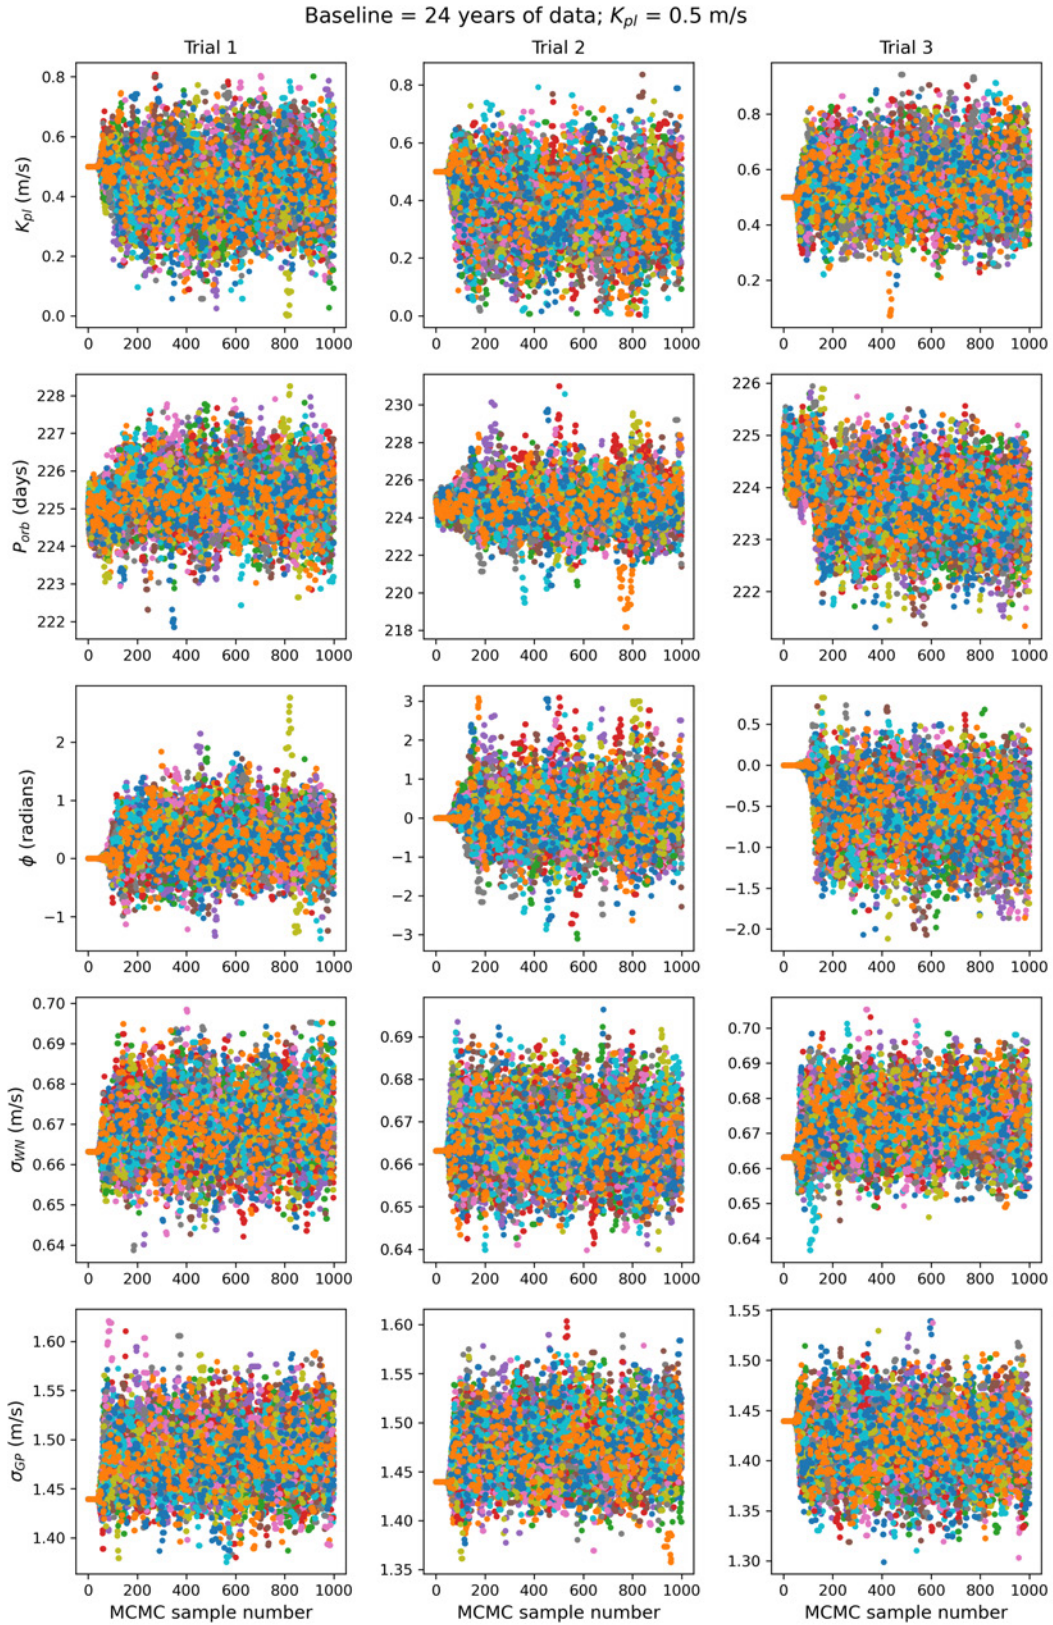


Figure A.10: MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.5 m/s. The data is synthesized with a baseline of 24 years.

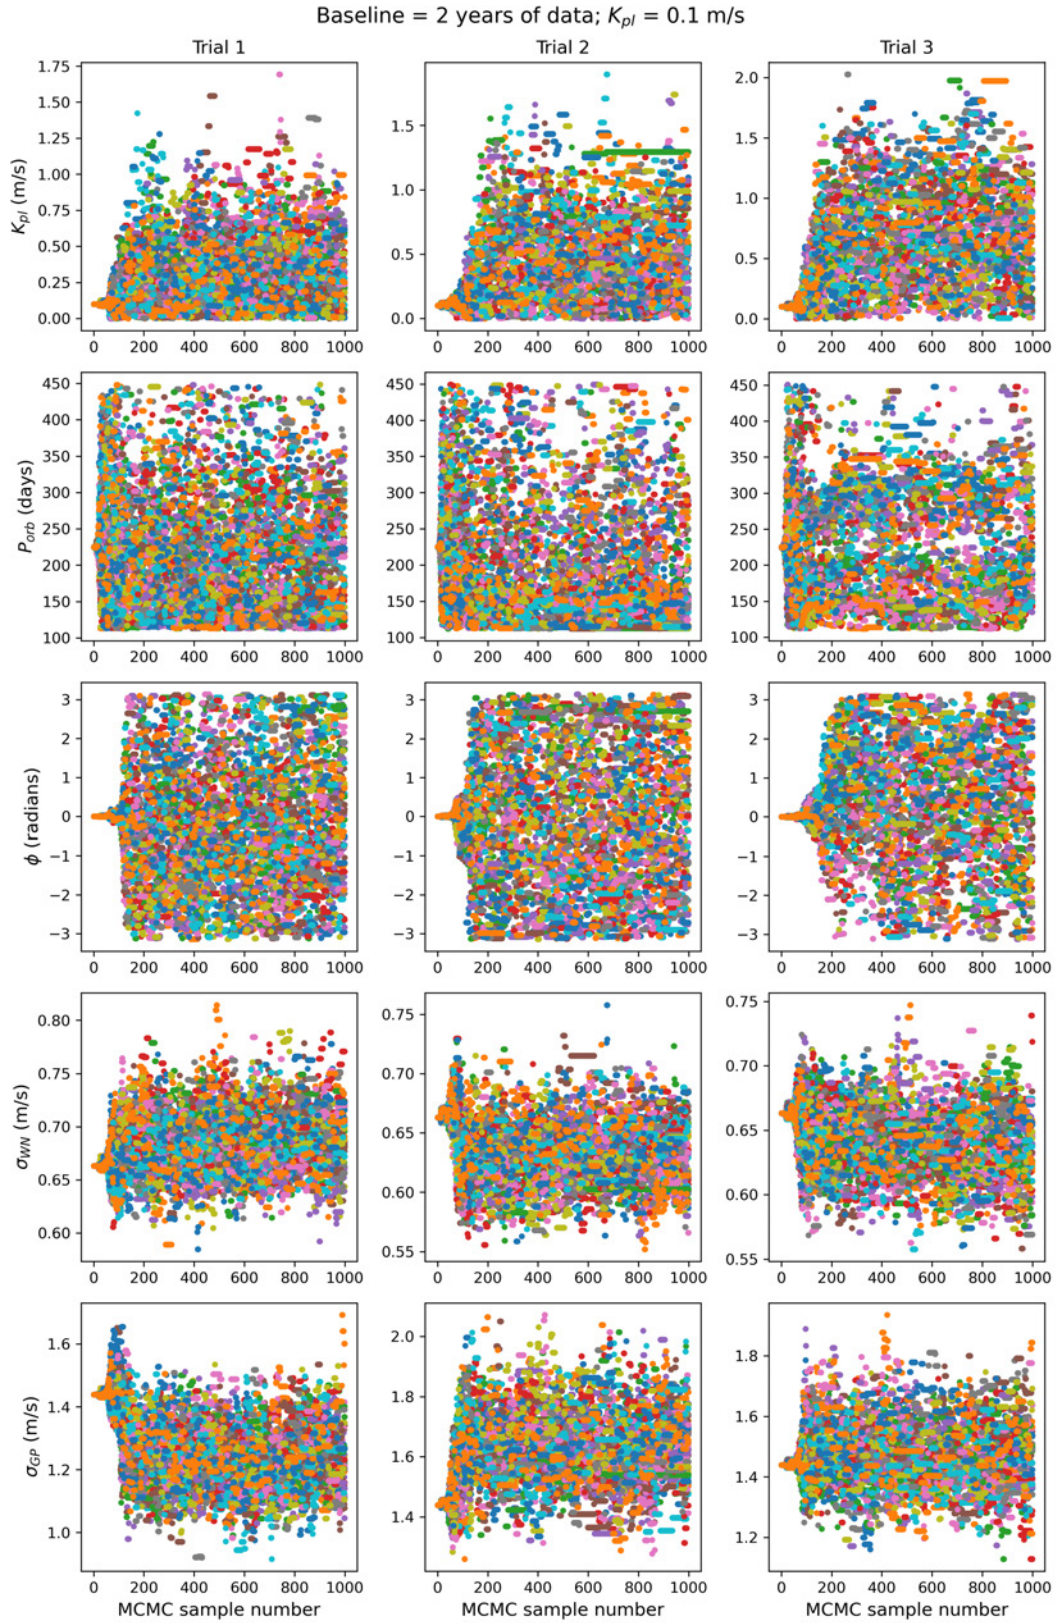


Figure A.11: MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.1 m/s. The data is synthesized with a baseline of 2 years.

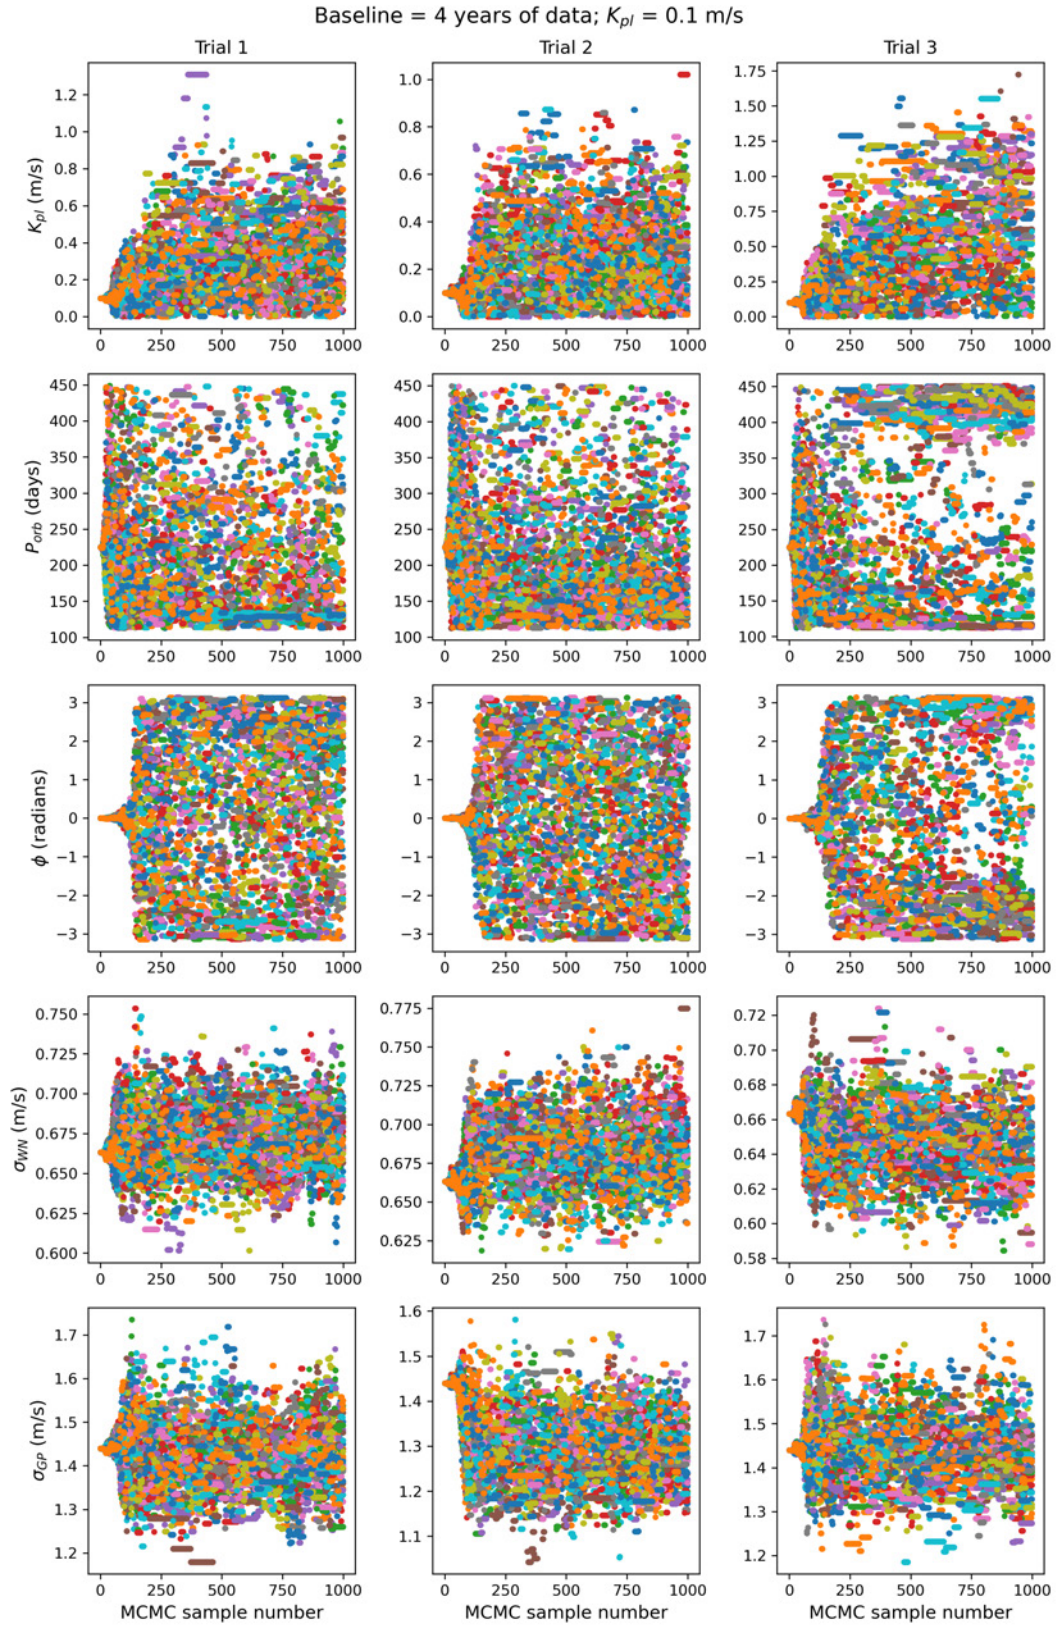


Figure A.12: MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.1 m/s. The data is synthesized with a baseline of 4 years.

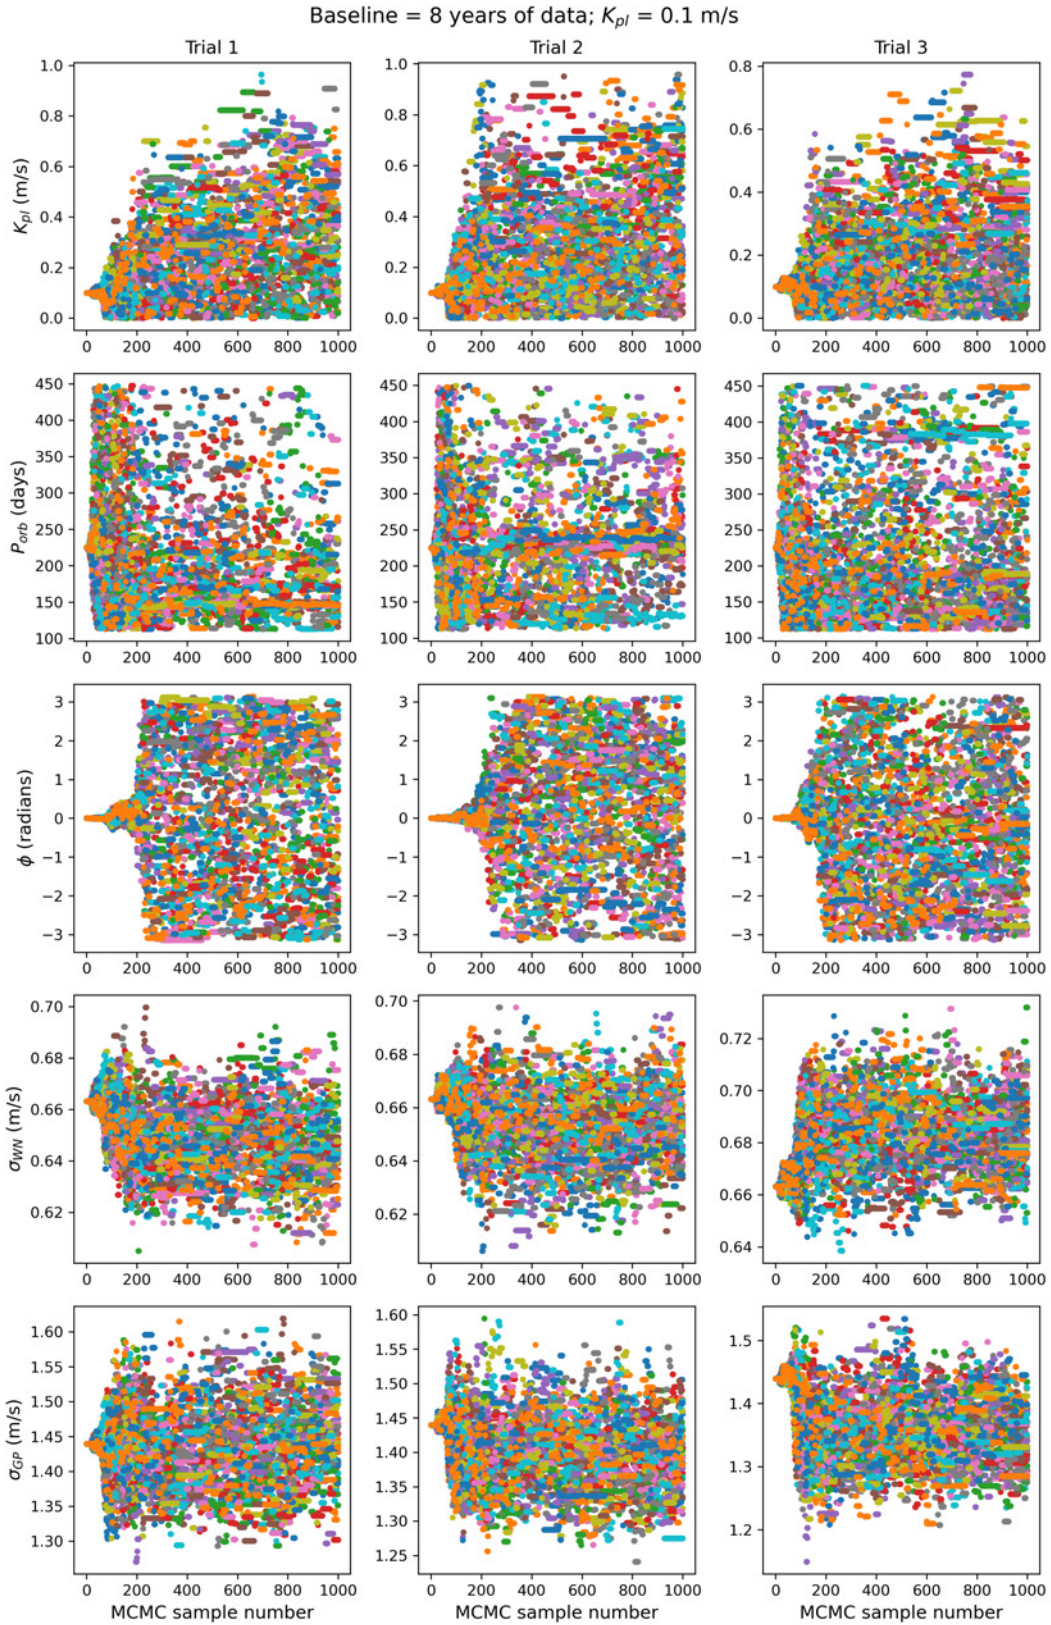


Figure A.13: MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.1 m/s. The data is synthesized with a baseline of 8 years.

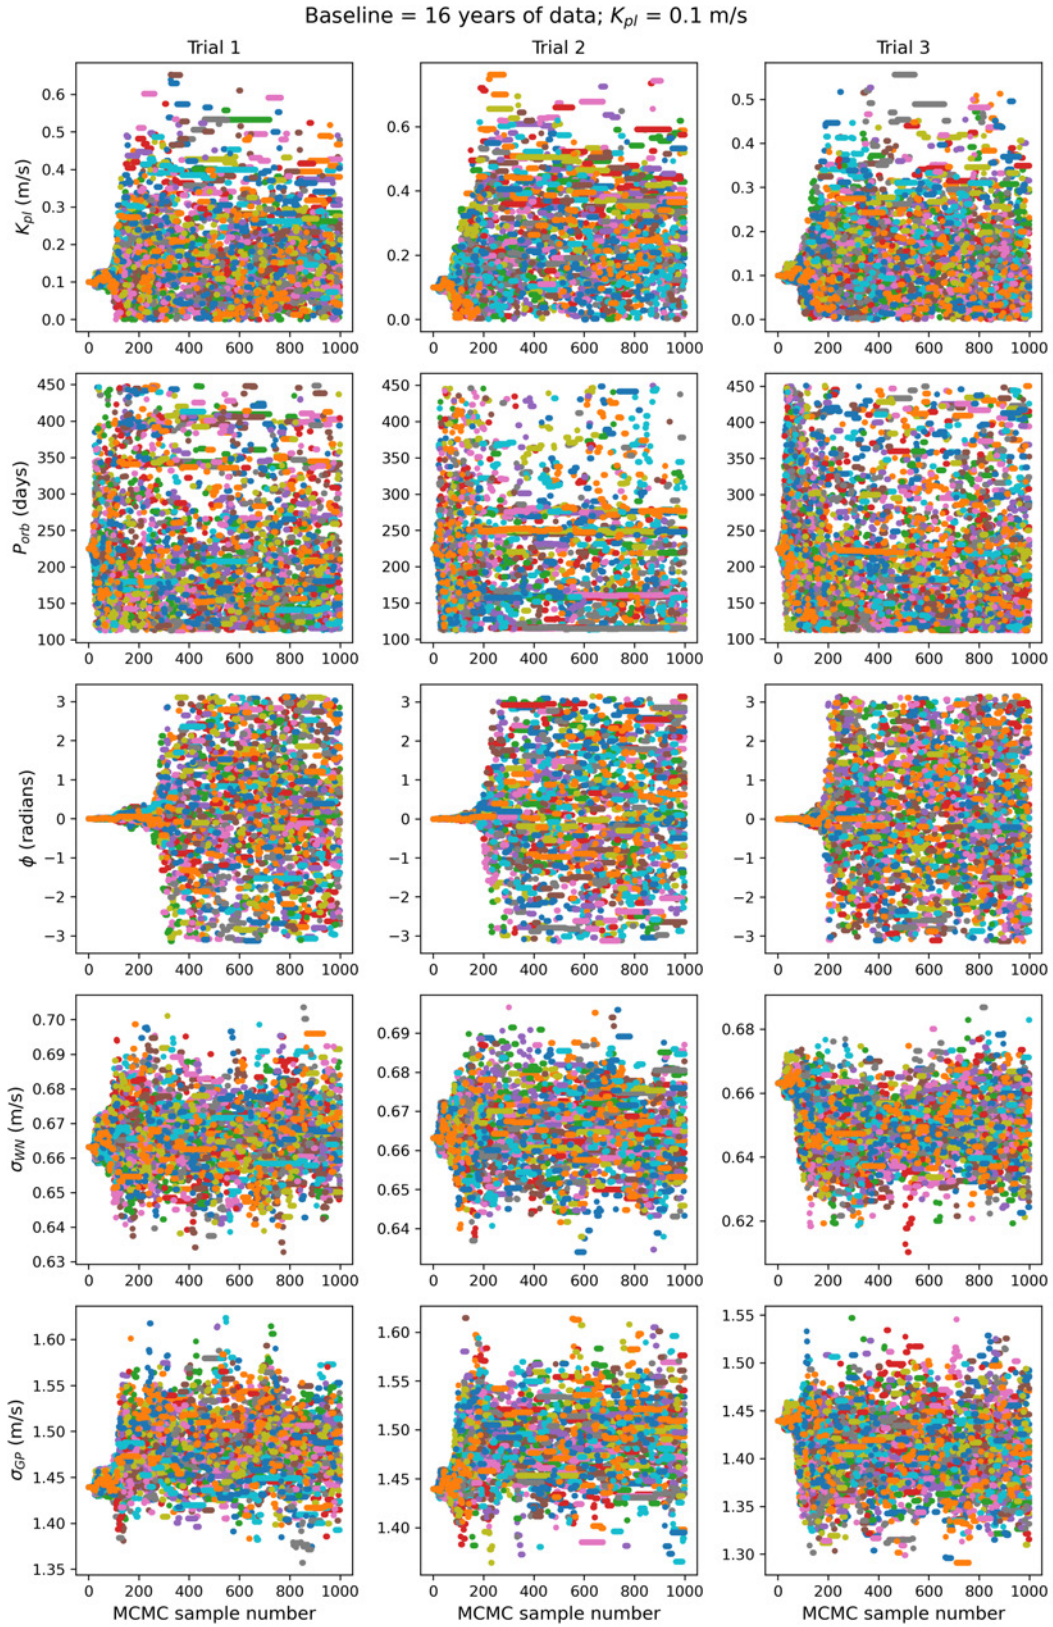


Figure A.14: MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.1 m/s. The data is synthesized with a baseline of 16 years.

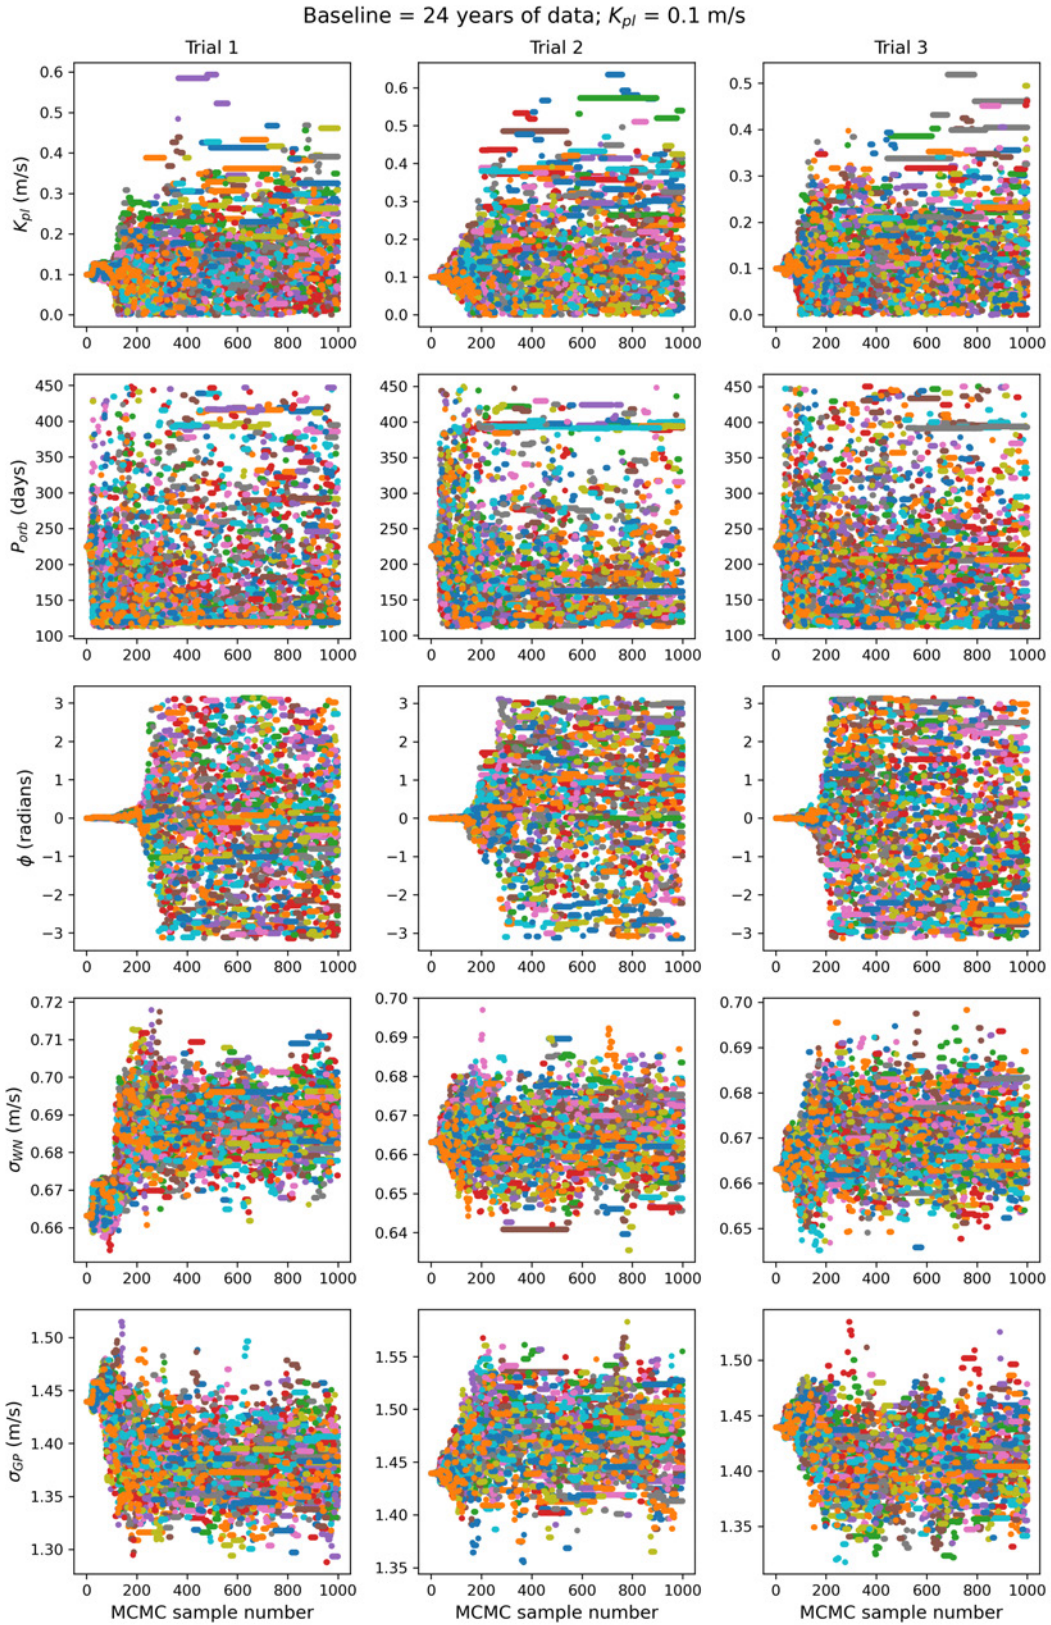


Figure A.15: MCMC samples for a synthetic 225 day planet with a semi-amplitude of 0.1 m/s. The data is synthesized with a baseline of 24 years.

B

Probing local dark matter density using stellar accelerations

In this appendix a summary is provided of the work in [157]. D. F. Phillips and R. L. Walsworth conceived the idea. M. Buschmann and B. R. Safdi performed the dark matter theory. A. Ravi created the simulation and analysis of future work. N. Langellier performed the analysis of the correlated noise dataset.

B.1 Introduction

Understanding the nature and distribution of dark matter is one of the most important issues in modern physics [158]. Current estimates of the local dark matter density rely on indirect measurements with large errors [159]. Here we propose a direct measure of the local dark matter density by observing the acceleration of clusters of stars in the Milky Way galaxy. In the absence of dark matter, the local acceleration of each star may be calculated as a function of its distance from the center of the Milky Way, r . Two stars with differing values of r will thus experience a difference in local radial acceleration. Shown in Figure B.1 is a diagram depicting the Earth located r_0 from the galactic center with a velocity v_0 and acceleration a_0 . A cluster of stars Δr towards the galactic




Figure B.1: Geometry for observing stellar accelerations in the Milky Way. The solar system is at a distance r_0 from the Galactic Center (origin), has a rotational velocity v_0 and feels an acceleration a_0 due to the Milky Way gravitational potential. Stars further inward feel a stronger acceleration. From Earth, we can observe the radial velocity of stars Δr away. By measuring small changes in these velocities over time, we directly determine stellar accelerations and hence the Milky Way gravitational potential. The diagram above is not to scale and angles are exaggerated for effect. Figure credit: A. Ravi.

center will feel a difference in acceleration Δa_r . If we choose $\Delta r = 3$ kpc, the acceleration difference is $\Delta a_r = 1.5 \times 10^{-8} \text{ cm/s}^2$, which over the course of 10 years corresponds to a change in velocity of about 5 cm/s. This number is conveniently close to the 10 cm/s target radial velocity (RV) semi-amplitude of an Earth like planet discussed in [chapter 2](#). Because 10 years is such a short fraction of the orbital period of a star about the galactic center, this change in velocity will appear as linear. Furthermore, these stars will experience the same noise properties also discussed in [chapter 2](#). Thus we can monitor the RV of a collection stars 3 kpc towards the galactic center over a decade and fit this time series with same Gaussian process (GP) model used in exoplanet searches. The only addition will be a linear drift term to measure the radial acceleration of each star. If the mean difference in acceleration is statistically significantly different from expected value in the absence of dark matter, we can obtain a direct measure of the local dark matter density. Details can be found in [\[157\]](#).

B.2 Results

In a series of numerical experiments, we attempt to recover an injected ~ 5 cm/s/decade acceleration signal from a large ensemble of synthesized v_{total} time series for N simulated primary stars.

First, to assess what is possible with present-day state-of-the-art instruments and analysis techniques, we set the Gaussian white noise standard deviation σ_{WN} to 60 cm/s and add correlated noise with amplitudes ranging from 50 cm/s to 250 cm/s on a given star (the range of noise levels observed in the Sun [160]). In order to fit the more complicated noise model, we employ a Gaussian process (GP) regression with a quasiperiodic kernel function [51, 72, 161]. The GP regression is also able to simultaneously fit the Keplerian components of the time series (i.e., those associated with planets and stellar companions). With the GP fit ($N = 75,044$ stars) we obtain a mean (standard error of the mean σ given in parentheses) of $1.42(0.24) \times 10^{-8}$ cm/s². In contrast, a simple linear fit ($N = 72,425$ stars) to the time series yields a mean acceleration of $2.00(0.58) \times 10^{-8}$ cm/s². Thus the GP fit reduces the uncertainty by more than a factor of 2 compared to the linear fit, recovers the injected stellar acceleration signal of 1.5×10^{-8} cm/s² within 1σ and is $\approx 6\sigma$ away from a null result. The probability density functions for the fitted accelerations are shown in Figure B.2(a).

We can vary the observed number of stars N to determine what is required for a statistically significant detection. For our purposes, a detection at the level $n\sigma$ is defined as the mean lying $n\sigma$ away from zero. Figure B.2(b) shows σ as a function of N . Unfortunately, nearly 2×10^4 stars are required for a 3σ detection ($\sigma \approx 0.5 \times 10^{-8}$ cm/s²). This is a prohibitively large sample size given realistic quantities of observation time. However, we reiterate that N represents the number of stars in an initial sample before target selection.

As in the case of Earth-like exoplanet detection, current telescope hardware and analysis techniques do not allow for the detection of differences in stellar accelerations suitable for dark matter density mapping. We find here as well that next generation telescopes in addition to better models of stellar activity will allow for these accelerations to be measured with a high degree of statistical

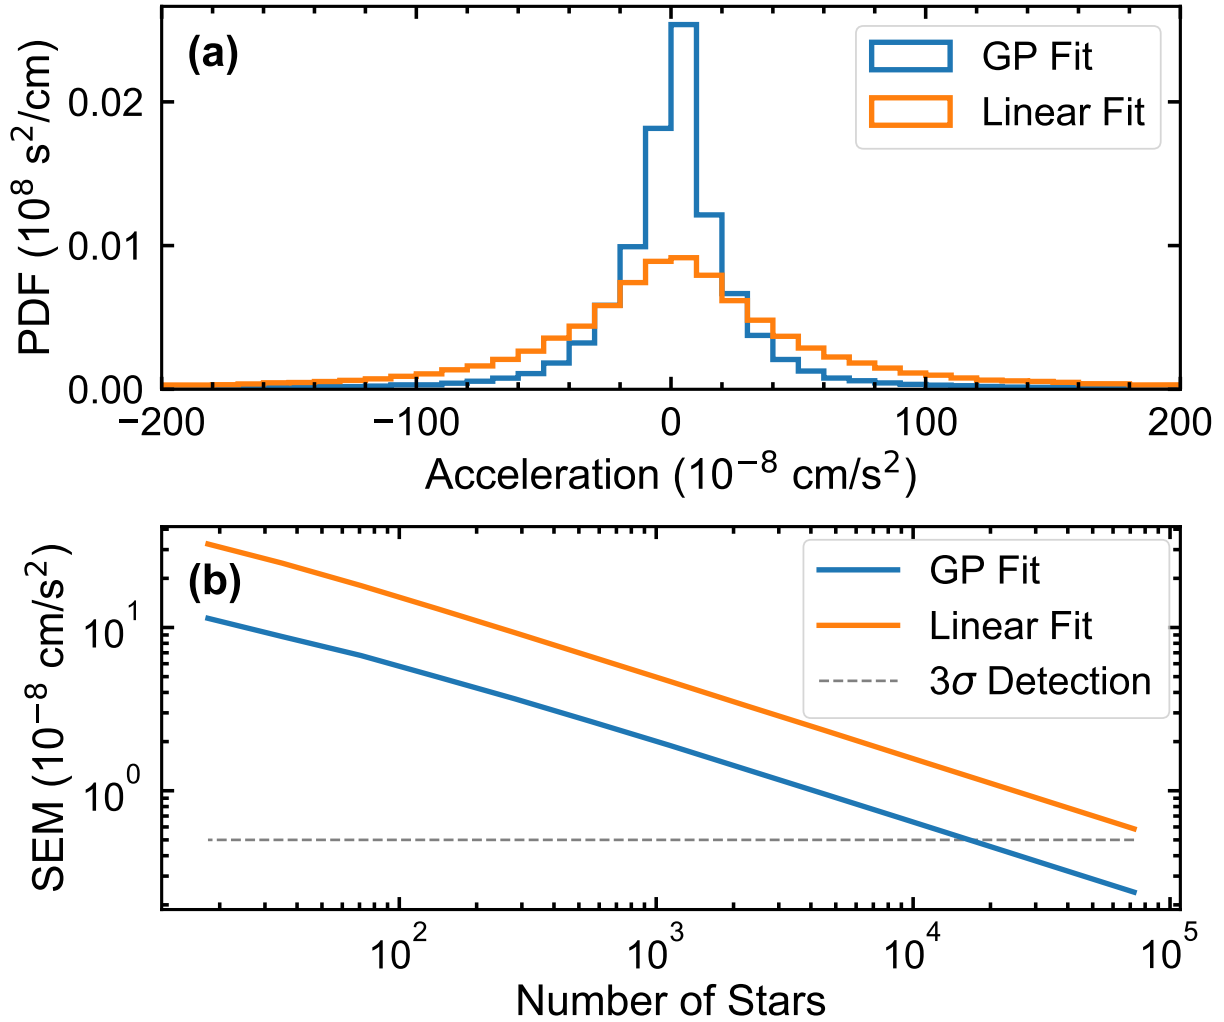


Figure B.2: (a) Probability density function (PDF) of fitted stellar accelerations for $\sim 7.5 \times 10^4$ stars using Gaussian processes (GP) (blue) and a simple linear fit (orange). 60 cm/s of white noise and quasiperiodic correlated noise ranging from 50 cm/s to 250 cm/s is added to each time series. (b) Standard error of the mean (SEM) vs. number of stars observed. A 3σ detection (dotted gray line) is obtained after about 2×10^4 stars for the GP fit.

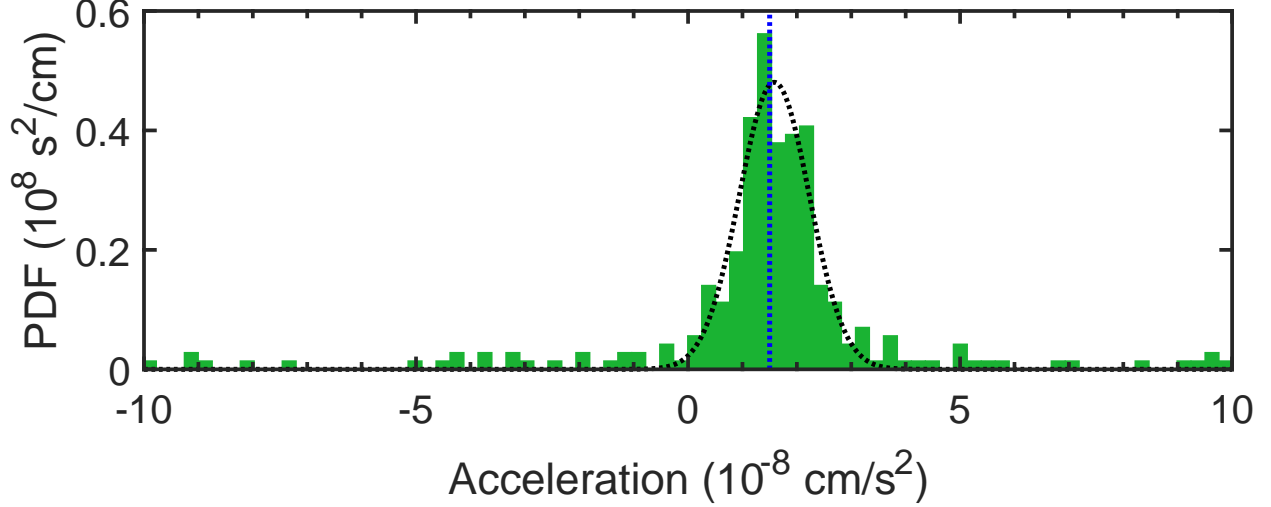


Figure B.3: Probability density function (PDF) of fitted stellar accelerations for 10^3 stars, using a simple linear fit and filtering using periodogram power. The only source of noise included is 10 cm/s white noise. The dotted black curve is a Gaussian fit to the histogram (a guide to the eye). The injected acceleration signal of $1.5 \times 10^{-8} \text{ cm/s}^2$ is shown as a blue dotted line. Figure credit: A. Ravi.

significance. Figure B.3 shows a histogram of linear fits to 1,000 RV time series with a white noise level of $\sigma_{\text{WN}} = 10 \text{ cm/s}$ and no correlated noise. Using these data, we obtain a mean acceleration of $1.46(0.21) \times 10^{-8} \text{ cm/s}^2$, which is a 7σ difference from a null result. This analysis shows promise for local dark matter density mapping with next generation telescopes.

B.3 Time series analysis using Gaussian processes

In order to test the feasibility of a stellar acceleration survey campaign with present-day technology and analysis techniques, realistic RV time series are generated with correlated noise and a Gaussian process (GP) regression is employed [51, 72, 161] to model these data. A GP assumes the data are normally distributed with mean $\boldsymbol{\mu}$ and covariance matrix \mathbf{K} . This allows for a simultaneous fit of any deterministic physics with a model of the correlated noise.

For this analysis, the mean function models both the stellar acceleration and M Keplerian signals

(stellar companions and planets):

$$\mu(t_i) = \Delta a_r t_i + v_{\text{off}} + \sum_{j=1}^M A_j \sin\left(\frac{2\pi t_i}{T_j} + \phi_j\right), \quad (\text{B.1})$$

where the relative stellar acceleration is Δa_r and the offset velocity is v_{off} . Each orbiting body is assumed to have zero eccentricity so A_j , T_j , and ϕ_j are the RV amplitude, orbital period, and orbital phase offset for the j^{th} orbiting body, respectively. The number of Keplerian signals, M , is chosen by assuming we would have knowledge of some (but not necessarily all) orbiting bodies for a given star. For our fits, we look at each orbiting body injected into the simulation and include it in Eq. (B.1) if the RV amplitude is greater than 1 m/s and the orbital period is less than 10^4 days. This corresponds to an optimistic cutoff where an orbiting body at the limit of state of the art techniques (~ 1 m/s) has an orbital period almost three times the observing campaign length. Such a body would have a change in RV over the observing campaign comparable to the injected instrumental white noise (~ 60 cm/s). Keeping all orbiting bodies that meet these criteria provide a best case scenario wherein the feasibility of this analysis is easily tested.

The covariance matrix \mathbf{K} is populated using the quasiperiodic kernel function given in Eq. (2.1) with an additive white noise term on the main diagonal:

$$K_{i,j}(t_i, t_j) = \sigma_{\text{QP}}^2 k_{\text{QP}}(t_i, t_j) + \sigma_{\text{WN}}^2 \delta(t_i, t_j) \quad (\text{B.2})$$

The GP model thus contains a total of $7 + 3M$ fit parameters,

$$\boldsymbol{\theta}_{\text{GP}} = \langle \Delta a_r, v_{\text{off}}, \{A_j\}, \{T_j\}, \{\phi_j\}, \sigma_{\text{QP}}, \tau, \eta, P_{\text{rot}}, \sigma_{\text{WN}} \rangle \quad (\text{B.3})$$

with a likelihood function given by

$$\mathcal{L}(\boldsymbol{\theta}_{\text{GP}} | \mathbf{v}, \mathbf{t}) = \frac{1}{\sqrt{|2\pi\mathbf{K}|}} e^{-\frac{1}{2}(\mathbf{v}-\boldsymbol{\mu})^T \mathbf{K}^{-1}(\mathbf{v}-\boldsymbol{\mu})}, \quad (\text{B.4})$$

where \mathbf{v} is a vector of the measured RVs at times \mathbf{t} , and $\boldsymbol{\mu}$ and \mathbf{K} are functions of \mathbf{t} and $\boldsymbol{\theta}_{\text{GP}}$ as described in Equations (B.1) and (B.2). The relative stellar acceleration, Δa_r , is given a Gaussian prior with a mean equal to the injected stellar acceleration signal and a standard deviation of 30 times the mean, i.e.,

$$P(\Delta a_r) \sim \mathcal{N}\left(1.5 \times 10^{-8} \text{ cm/s}^2, 45 \times 10^{-8} \text{ cm/s}^2\right). \quad (\text{B.5})$$

This prior is chosen to include the expected value of the relative stellar acceleration and reasonable deviations about this value while rejecting accelerations from long-period orbiting bodies that are degenerate with the stellar acceleration signal on the timescale of a decade. The remaining model parameters are allowed to vary freely.

The maximum a posteriori is taken as the parameter estimate, $\hat{\boldsymbol{\theta}}_{\text{GP}}$. Minimization of the negative log posterior yields the desired result:

$$\hat{\boldsymbol{\theta}}_{\text{GP}} = \text{argmin}_{\boldsymbol{\theta}_{\text{GP}}} [-\ln L(\boldsymbol{\theta}_{\text{GP}}|\mathbf{v}, \mathbf{t}) - \ln P(\Delta a_r)]. \quad (\text{B.6})$$

Minimization is performed using the quasi-Newton method of Broyden, Fletcher, Goldfarb, and Shanno [162]. Initial guesses for the model parameters, $\boldsymbol{\theta}_{\text{GP},0}$, were chosen assuming we would have knowledge of their true values, $\boldsymbol{\theta}_{\text{GP, true}}$. The values were normally distributed about their true value with a standard deviation of about 3% of the mean, i.e.,

$$P(\boldsymbol{\theta}_{\text{GP},0}) \sim \mathcal{N}\left(\boldsymbol{\theta}_{\text{GP, true}}, \frac{1}{30}\boldsymbol{\theta}_{\text{GP, true}}\right). \quad (\text{B.7})$$

The Python [76] packages NumPy [77], SciPy [78], and George [80] were used for this analysis.

Bibliography

- [1] B. P. Abbott, R. Abbott, T. D. Abbott, M. R. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. X. Adhikari, *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **116**, 061102 (2016).
- [2] B. Abbott, R. Abbott, T. Abbott, M. Abernathy, F. Acernese, K. Ackley, C. Adams, T. Adams, P. Addesso, R. Adhikari, *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), *Phys. Rev. Lett.* **116**, 241103 (2016).
- [3] B. Abbott, R. Abbott, R. Adhikari, P. Ajith, B. Allen, G. Allen, R. Amin, S. Anderson, W. Anderson, M. Arain, *et al.*, *Reports on Progress in Physics* **72**, 076901 (2009).
- [4] K. Ishidoshiro, M. Ando, A. Takamori, H. Takahashi, K. Okada, N. Matsumoto, W. Kokuyama, N. Kanda, Y. Aso, and K. Tsubono, *Phys. Rev. Lett.* **106**, 161101 (2011).
- [5] M. Coughlin and J. Harms, *Phys. Rev. Lett.* **112**, 101102 (2014).
- [6] S. Detweiler, *The Astrophysical Journal* **234**, 1100 (1979).
- [7] S. Aoyama, R. Tazai, and K. Ichiki, *Phys. Rev. D* **89**, 067101 (2014).
- [8] B. Bertotti, R. Ambrosini, J. Armstrong, S. Asmar, G. Comoretto, G. Giampieri, L. Iess, Y. Koyama, A. Messeri, A. Vecchio, *et al.*, *Astronomy and Astrophysics* **296**, 13 (1995).
- [9] A. Konopliv, S. Asmar, E. Carranza, W. Sjogren, and D. Yuan, *Icarus* **150**, 1 (2001).
- [10] J. W. Armstrong, *Living Rev. Relativity* **9**, E2 (2006).
- [11] B. J. Bloom, T. L. Nicholson, J. R. Williams, S. L. Campbell, M. Bishof, X. Zhang, W. Zhang, S. L. Bromley, and J. Ye, *Nature* **506**, 71 (2014).
- [12] N. Hinkley, J. A. Sherman, N. B. Phillips, M. Schioppo, N. D. Lemke, K. Beloy, M. Pizzocaro, C. W. Oates, and A. D. Ludlow, *Science* **341**, 1215 (2013).
- [13] I. Ushijima, M. Takamoto, M. Das, T. Ohkubo, and H. Katori, *Nature Photonics* **9**, 185 (2015).

- [14] T. Nicholson, S. Campbell, R. Hutson, G. Marti, B. Bloom, R. McNally, W. Zhang, M. Barrett, M. Safronova, G. Strouse, *et al.*, *Nature Comm.* **6** (2015).
- [15] N. Huntemann, C. Sanner, B. Lipphardt, C. Tamm, and E. Peik, *Phys. Rev. Lett.* **116**, 063001 (2016).
- [16] A. Loeb and D. Maoz, arXiv preprint arXiv:1501.00996 (2015).
- [17] A. Vutha, *New Journal of Physics* **17**, 063030 (2015).
- [18] K. Danzmann, T. Prince, P. Binétruy, P. Bender, S. Buchman, J. Centrella, M. Cerdonio, N. Cornish, A. Cruise, C. Cutler, *et al.*, *Assessment Study Report ESA/SRE* **3**, 2 (2011).
- [19] P. Amaro-Seoane, S. Aoudia, S. Babak, P. Binétruy, E. Berti, A. Bohe, C. Caprini, M. Colpi, N. J. Cornish, K. Danzmann, *et al.*, *Classical and Quantum Gravity* **29**, 124016 (2012).
- [20] P. W. Graham, J. M. Hogan, M. A. Kasevich, and S. Rajendran, *Phys. Rev. Lett.* **110**, 171102 (2013).
- [21] J. M. Hogan and M. A. Kasevich, *Phys. Rev. A* **94**, 033632 (2016).
- [22] N. Yu and M. Tinto, *General Relativity and Gravitation* **43**, 1943 (2011).
- [23] S.-W. Chiow, J. Williams, and N. Yu, *Phys. Rev. A* **92**, 063613 (2015).
- [24] W. Chaibi, R. Geiger, B. Canuel, A. Bertoldi, A. Landragin, and P. Bouyer, *Phys. Rev. D* **93**, 021101 (2016).
- [25] G. Congedo, R. Dolesi, M. Hueller, S. Vitale, and W. J. Weber, **88**, 082003 (2013), [arXiv:1306.1422 \[gr-qc\]](#) .
- [26] J. Cordes and R. Shannon, arXiv preprint arXiv:1010.3785 (2010).
- [27] C. W. Chou, D. B. Hume, J. C. J. Koelemeij, D. J. Wineland, and T. Rosenband, *Phys. Rev. Lett.* **104**, 070802 (2010).
- [28] M. Takamoto, T. Takano, and H. Katori, *Nature Photonics* **5**, 288 (2011).
- [29] T. Takano, M. Takamoto, I. Ushijima, N. Ohmae, T. Akatsuka, A. Yamaguchi, Y. Kuroishi, H. Munekane, B. Miyahara, and H. Katori, *Nature Photonics* (2016).
- [30] T. Nicholson, M. Martin, J. Williams, B. Bloom, M. Bishof, M. Swallows, S. Campbell, and J. Ye, *Phys. Rev. Lett.* **109**, 230801 (2012).
- [31] M. Bishof, X. Zhang, M. J. Martin, and J. Ye, *Phys. Rev. Lett.* **111**, 093604 (2013).
- [32] J. Taylor, P. Cappellaro, L. Childress, L. Jiang, D. Budker, P. Hemmer, A. Yacoby, R. Walsworth, and M. Lukin, *Nature Physics* **4**, 810 (2008).
- [33] I. D. Leroux, M. H. Schleier-Smith, and V. Vuletić, *Phys. Rev. Lett.* **104**, 073602 (2010).

- [34] J. G. Bohnet, K. C. Cox, M. A. Norcia, J. M. Weiner, Z. Chen, and J. K. Thompson, *Nature Photonics* **8**, 731 (2014).
- [35] M. Kessler, E., P. Komar, M. Bishof, L. Jiang, S. Sorensen, A., J. Ye, and D. Lukin, M., *Phys. Rev. Lett.* **112**, 190403 (2014).
- [36] A. Derevianko and M. Pospelov, *Nature Physics* **10**, 933 (2014).
- [37] A. Arvanitaki, J. Huang, and K. Van Tilburg, *Phys. Rev. D* **91**, 015015 (2015).
- [38] A. Arvanitaki, P. W. Graham, J. M. Hogan, S. Rajendran, and K. Van Tilburg, arXiv preprint arXiv:1606.04541 (2016).
- [39] C. Delaunay and Y. Soreq, arXiv preprint arXiv:1602.04838 (2016).
- [40] S. Origlia, S. Schiller, M. Pramod, L. Smith, Y. Singh, W. He, S. Viswam, D. Świerad, J. Hughes, K. Bongs, *et al.*, in *Proc. SPIE 9900, Quantum Optics* (2016) p. 990003.
- [41] T. Kessler, C. Hagemann, C. Grebing, T. Legero, U. Sterr, F. Riehle, M. J. Martin, L. Chen, and J. Ye, *Nature Photonics* **6**, 687 (2012).
- [42] P. McNamara, S. Vitale, K. Danzmann, L. P. S. W. Team, *et al.*, *Classical and Quantum Gravity* **25**, 114034 (2008).
- [43] M. Armano, H. Audley, G. Auger, J. Baird, M. Bassan, P. Binetruy, M. Born, D. Bortoluzzi, N. Brandt, M. Caleno, *et al.*, *Phys. Rev. Lett.* **116**, 231101 (2016).
- [44] S. M. Foreman, A. D. Ludlow, M. H. de Miranda, J. E. Stalnaker, S. A. Diddams, and J. Ye, *Phys. Rev. Lett.* **99**, 153601 (2007).
- [45] G. Heinzel, V. Wand, A. Garcia, O. Jennrich, C. Braxmaier, D. Robertson, K. Middleton, D. Hoyland, A. Rüdiger, R. Schilling, *et al.*, *Classical and Quantum Gravity* **21**, S581 (2004).
- [46] S. L. Larson, W. A. Hiscock, and R. W. Hellings, *Phys. Rev. D* **62**, 062001 (2000).
- [47] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P. O. Schmidt, *Reviews of Modern Physics* **87**, 637 (2015).
- [48] F. B. Estabrook and H. D. Wahlquist, *General Relativity and Gravitation* **6**, 439 (1975).
- [49] Cegla, H. M., *Geosciences* **9**, 114 (2019).
- [50] Fischer, D. A., G. Anglada-Escude, P. Arriagada, R. V. Baluev, J. L. Bean, F. Bouchy, L. A. Buchhave, T. Carroll, A. Chakraborty, J. R. Crepp, R. I. Dawson, S. A. Diddams, X. Dumusque, J. D. Eastman, M. Endl, P. Figueira, E. B. Ford, D. Foreman-Mackey, P. Fournier, G. Fűrész, B. S. Gaudi, P. C. Gregory, F. Grundahl, A. P. Hatzes, G. Hébrard, E. Herrero, D. W. Hogg, A. W. Howard, J. A. Johnson, P. Jordan, C. A. Jurgenson, D. W. Latham, G. Laughlin, T. J. Loredo, C. Lovis, S. Mahadevan, T. M. McCracken, F. Pepe, M. Perez, D. F. Phillips, P. P. Plavchan, L. Prato, A. Quirrenbach, A. Reiners, P. Robertson, N. C. Santos, D. Sawyer, D. Segransan, A. Sozzetti, T. Steinmetz, A. Szentgyorgyi, S. Udry, J. A. Valenti, S. X. Wang, R. A. Wittenmyer, and J. T. Wright, *PASP* **128**, 066001 (2016), arXiv:1602.07939

[astro-ph.IM] .

- [51] Haywood, R. D., A. Collier Cameron, D. Queloz, S. C. C. Barros, M. Deleuil, R. Fares, M. Gillon, A. F. Lanza, C. Lovis, C. Moutou, F. Pepe, D. Pollacco, A. Santerne, D. Ségransan, and Y. C. Unruh, *MNRAS* **443**, 2517 (2014), [arXiv:1407.1044 \[astro-ph.EP\]](#) .
- [52] Robertson, P., S. Mahadevan, M. Endl, and A. Roy, *Science* **345**, 440 (2014), <http://science.sciencemag.org/content/345/6195/440.full.pdf> .
- [53] Rajpaul, V., S. Aigrain, and S. Roberts, *MNRAS* **456**, L6 (2016), [arXiv:1510.05598 \[astro-ph.EP\]](#) .
- [54] Kjeldsen, H. and Bedding, T. R., *A&A* **293**, 87 (1995), [astro-ph/9403015](#) .
- [55] Giles, H., A. C., A. Collier Cameron, and R. D. Haywood, *MNRAS* **472**, 1618 (2017), <http://oup.prod.sis.lan/mnras/article-pdf/472/2/1618/19914219/stx1931.pdf> .
- [56] Meunier, N., A.-M. Lagrange, and M. Desort, *A&A* **519**, A66 (2010), [arXiv:1005.4764 \[astro-ph.SR\]](#) .
- [57] Dumusque, X., I. Boisse, and N. C. Santos, *ApJ* **796**, 132 (2014), [arXiv:1409.3594 \[astro-ph.SR\]](#) .
- [58] Hall, R. D., S. J. Thompson, W. Handley, and D. Queloz, *Monthly Notices of the Royal Astronomical Society* **479**, 2968 (2018), <https://academic.oup.com/mnras/article-pdf/479/3/2968/25152756/sty1464.pdf> .
- [59] Haywood, R. D., T. W. Milbourne, S. H. Saar, A. Mortier, D. Phillips, D. Charbonneau, A. C. Cameron, H. M. Cegla, N. Meunier, and M. L. P. III, “Unsigned magnetic flux as a proxy for radial-velocity variations in sun-like stars,” (2020), [arXiv:2005.13386 \[astro-ph.SR\]](#) .
- [60] Dumusque, X., Borsa, F., Damasso, M., Díaz, R. F., Gregory, P. C., Hara, N. C., Hatzes, A., Rajpaul, V., Tuomi, M., Aigrain, S., Anglada-Escudé, G., Bonomo, A. S., Boué, G., Dauvergne, F., Frustagli, G., Giacobbe, P., Haywood, R. D., Jones, H. R. A., Laskar, J., Pinamonti, M., Poretti, E., Rainer, M., Ségransan, D., Sozzetti, A., and Udry, S., *A&A* **598**, A133 (2017).
- [61] Damasso, M., Zeng, L., Malavolta, L., Mayo, A., Sozzetti, A., Mortier, A., Buchhave, L. A., Vanderburg, A., Lopez-Morales, M., Bonomo, A. S., Cameron, A. C., Coffinet, A., Figueira, P., Latham, D. W., Mayor, M., Molinari, E., Pepe, F., Phillips, D. F., Poretti, E., Rice, K., Udry, S., and Watson, C. A., *A&A* **624**, A38 (2019).
- [62] Cosentino, R., C. Lovis, F. Pepe, A. C. Cameron, D. W. Latham, E. Molinari, S. Udry, N. Bezawada, N. Buchschacher, P. Figueira, M. Fleury, A. Ghedina, A. G. Glenday, M. Gonzalez, J. Guerra, D. Henry, I. Hughes, C. Maire, F. Motalebi, and D. F. Phillips, in *Ground-based and Airborne Instrumentation for Astronomy V*, Vol. 9147, edited by S. K. Ramsay, I. S. McLean, and H. Takami, International Society for Optics and Photonics (SPIE, 2014) pp. 2658 – 2669.
- [63] Dumusque, X., A. Glenday, D. F. Phillips, N. Buchschacher, A. Collier Cameron, M. Ceconi,

- D. Charbonneau, R. Cosentino, A. Ghedina, D. W. Latham, C.-H. Li, M. Lodi, C. Lovis, E. Molinari, F. Pepe, S. Udry, D. Sasselov, A. Szentgyorgyi, and R. Walsworth, *ApJl* **814**, L21 (2015), [arXiv:1511.02267 \[astro-ph.EP\]](https://arxiv.org/abs/1511.02267) .
- [64] Phillips, D. F., A. G. Glenday, X. Dumusque, N. Buchschacher, A. C. Cameron, M. Ceconi, D. Charbonneau, R. Cosentino, A. Ghedina, R. Haywood, D. W. Latham, C.-H. Li, M. Lodi, C. Lovis, E. Molinari, F. Pepe, D. Sasselov, A. Szentgyorgyi, S. Udry, and R. L. Walsworth, *Proc.SPIE* **9912**, 9912 (2016).
- [65] Baranne, A., D. Queloz, M. Mayor, G. Adrianzyk, G. Knispel, D. Kohler, D. Lacroix, J.-P. Meunier, G. Rimbaud, and A. Vin, **119**, 373 (1996).
- [66] Sosnowska, D., M. Lodi, X. Gao, N. Buchschacher, A. Vick, J. Guerra, M. Gonzalez, D. Kelly, C. Lovis, F. Pepe, E. Molinari, A. C. Cameron, D. Latham, and S. Udry, *Proc.SPIE* **8451**, 8451 (2012).
- [67] Wilson, O. C., *Astrophysical Journal* **153**, 221 (1968).
- [68] Linksy, J. L. and Avrett, E. H., *Publications of the Astronomical Society of the Pacific* **82**, 169 (1970).
- [69] Collier Cameron, A, A. Mortier, D. Phillips, X. Dumusque, R. D. Haywood, N. Langellier, C. A. Watson, H. M. Cegla, J. Costes, D. Charbonneau, A. Coffinet, D. W. Latham, M. Lopez-Morales, L. Malavolta, J. Maldonado, G. Micela, T. Milbourne, E. Molinari, S. H. Saar, S. Thompson, N. Buchschacher, M. Ceconi, R. Cosentino, A. Ghedina, A. Glenday, M. Gonzalez, C.-H. Li, M. Lodi, C. Lovis, F. Pepe, E. Poretti, K. Rice, D. Sasselov, A. Sozzetti, A. Szentgyorgyi, S. Udry, and R. Walsworth, *Monthly Notices of the Royal Astronomical Society* **487**, 1082 (2019), <http://oup.prod.sis.lan/mnras/article-pdf/487/1/1082/28753321/stz1215.pdf> .
- [70] Isaacson, H. and Fischer, D., *The Astrophysical Journal* **725**, 875 (2010).
- [71] Motalebi, F., Udry, S., Gillon, M., Lovis, C., Ségransan, D., Buchhave, L. A., Demory, B. O., Malavolta, L., Dressing, C. D., Sasselov, D., Rice, K., Charbonneau, D., Collier Cameron, A., Latham, D., Molinari, E., Pepe, F., Affer, L., Bonomo, A. S., Cosentino, R., Dumusque, X., Figueira, P., Fiorenzano, A. F. M., Gettel, S., Harutyunyan, A., Haywood, R. D., Johnson, J., Lopez, E., Lopez-Morales, M., Mayor, M., Micela, G., Mortier, A., Nascimbeni, V., Philips, D., Piotto, G., Pollacco, D., Queloz, D., Sozzetti, A., Vanderburg, A., and Watson, C. A., *A&A* **584**, A72 (2015).
- [72] Rasmussen, C. E. and Williams, C. K. I., *Gaussian Processes for Machine Learning*, 2nd ed. (MIT Press, 2006).
- [73] Rajpaul, V., S. Aigrain, M. A. Osborne, S. Reece, and S. Roberts, *Monthly Notices of the Royal Astronomical Society* **452**, 2269 (2015), <https://academic.oup.com/mnras/article-pdf/452/3/2269/4912584/stv1428.pdf> .
- [74] Faria, J. P., Haywood, R. D., Brewer, B. J., Figueira, P., Oshagh, M., Santerne, A., and Santos, N. C., *A&A* **588**, A31 (2016).
- [75] Grunblatt, S. K., A. W. Howard, and R. D. Haywood, *The Astrophysical Journal* **808**, 127

- (2015).
- [76] van Rossum, G., *Python tutorial*, Tech. Rep. CS-R9526 (Centrum voor Wiskunde en Informatica (CWI), Amsterdam, 1995).
- [77] Oliphant, T., “NumPy: A guide to NumPy,” USA: Trelgol Publishing (2006).
- [78] Jones, E., T. Oliphant, P. Peterson, *et al.*, “SciPy: Open source scientific tools for Python,” (2001).
- [79] Foreman-Mackey, D., D. W. Hogg, D. Lang, and J. Goodman, **125**, 306 (2013), [arXiv:1202.3665 \[astro-ph.IM\]](#) .
- [80] Ambikasaran, S., D. Foreman-Mackey, L. Greengard, D. W. Hogg, and M. O’Neil, “Fast direct methods for gaussian processes,” (2014), [arXiv:1403.6015](#) .
- [81] Student, *Biometrika* **6**, 302 (1908).
- [82] Milbourne, T. W., R. D. Haywood, D. F. Phillips, S. H. Saar, H. M. Cegla, A. C. Cameron, J. Costes, X. Dumusque, N. Langellier, D. W. Latham, J. Maldonado, L. Malavolta, A. Mortier, M. L. P. III, S. Thompson, C. A. Watson, F. Bouchy, N. Buchschacher, M. Ceconi, D. Charbonneau, R. Cosentino, A. Ghedina, A. G. Glenday, M. Gonzalez, C.-H. Li, M. Lodi, M. López-Morales, C. Lovis, M. Mayor, G. Micela, E. Molinari, F. Pepe, G. Pitozzo, K. Rice, D. Sasselov, D. Ségransan, A. Sozzetti, A. Szentgyorgyi, S. Udry, and R. L. Walsworth, *The Astrophysical Journal* **874**, 107 (2019).
- [83] Miklos, M., T. W. Milbourne, R. D. Haywood, D. F. Phillips, S. H. Saar, N. Meunier, H. M. Cegla, X. Dumusque, N. Langellier, J. Maldonado, L. Malavolta, A. Mortier, S. Thompson, C. A. Watson, M. Ceconi, R. Cosentino, A. Ghedina, C.-H. Li, M. López-Morales, E. Molinari, E. Poretti, D. Sasselov, A. Sozzetti, and R. L. Walsworth, *The Astrophysical Journal* **888**, 117 (2020).
- [84] Perryman, M., *The Exoplanet Handbook* (Cambridge University Press, 2011).
- [85] Nava, C., M. López-Morales, R. D. Haywood, and H. A. C. Giles, *The Astronomical Journal* **159**, 23 (2019).
- [86] Scott, D. W., *Multivariate Density Estimation: Theory, Practice, and Visualization*, 2nd ed. (John Wiley and Sons, Chicester, New York, 2015).
- [87] Cloutier, R., R. Doyon, F. Bouchy, and G. Hébrard, *The Astronomical Journal* **156**, 82 (2018).
- [88] Aigrain, S., F. Pont, and S. Zucker, *Monthly Notices of the Royal Astronomical Society* **419**, 3147 (2012), <https://academic.oup.com/mnras/article-pdf/419/4/3147/9506243/mnras0419-3147.pdf> .
- [89] Dumusque, X., *A&A* **620**, A47 (2018).
- [90] Cretignier, M., Francfort, J., Dumusque, X., Allart, R., and Pepe, F., “Rassine: Interac-

- tive tool for normalising stellar spectra i. description and performance of the code,” (2020), [arXiv:2006.13098](https://arxiv.org/abs/2006.13098) .
- [91] Meunier, N., Lagrange, A.-M., and Borgniet, S., *A&A* **607**, A6 (2017).
- [92] Wright, J. T. and Robertson, P., *Research Notes of the American Astronomical Society* **1**, 51 (2017), [arXiv:1801.05383 \[astro-ph.IM\]](https://arxiv.org/abs/1801.05383) .
- [93] Pepe, F. A., S. Cristiani, R. Rebolo Lopez, N. C. Santos, A. Amorim, G. Avila, W. Benz, P. Bonifacio, A. Cabral, P. Carvas, R. Cirami, J. Coelho, M. Comari, I. Coretti, V. De Caprio, H. Dekker, B. Delabre, P. Di Marcantonio, V. D’Odorico, M. Fleury, R. García, J. M. Herberos Linares, I. Hughes, O. Iwert, J. Lima, J.-L. Lizon, G. Lo Curto, C. Lovis, A. Manescau, C. Martins, D. Mégevand, A. Moitinho, P. Molaro, M. Monteiro, M. Monteiro, L. Pasquini, C. Mordasini, D. Queloz, J. L. Rasilla, J. M. Rebordão, S. Santana Tschudi, P. Santin, D. Sosnowska, P. Spanò, F. Tenegi, S. Udry, E. Vanzella, M. Viel, M. R. Zapatero Osorio, and F. Zerbi, in , *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7735 (2010) p. 77350F.
- [94] Allen, L. E., J. Wright, J. Rajagopal, F. Santoro, M. Liang, E. Timmerman, R. Christensen, E. Hunting, M. Wolf, K. Jaehrig, J. Percival, M. Smith, S. Mahadevan, F. Hearty, C. Bender, C. Blake, S. E. Logsdon, R. Akeson, R. W. Capps, J. Callas, P. A. Willems, M. McElwain, M. W. McElwain, F. Basten, A. Monson, G. Stefansson, L. Ramsey, J. Ninan, S. Blakeslee, K. Kaplan, S. Halverson, A. Roy, R. Terrien, P. Robertson, C. Schwab, M. Rud, and S. Kandodia, in *American Astronomical Society Meeting Abstracts #231*, *American Astronomical Society Meeting Abstracts*, Vol. 231 (2018) p. 246.08.
- [95] Jurgenson, C., D. Fischer, T. McCracken, D. Sawyer, A. Szymkowiak, A. Davis, G. Muller, and F. Santoro, in *Ground-based and Airborne Instrumentation for Astronomy VI*, Vol. 9908, edited by C. J. Evans, L. Simard, and H. Takami, *International Society for Optics and Photonics (SPIE)*, (2016) pp. 2051 – 2070.
- [96] Thompson, S. J., D. Queloz, I. Baraffe, M. Brake, A. Dolgoplov, M. Fisher, M. Fleury, J. Geelhoed, R. Hall, J. I. G. Hernández, R. ter Horst, J. Kragt, R. Navarro, T. Naylor, F. Pepe, N. Piskunov, R. Rebolo, L. Sander, D. SÁ©gransan, E. Seneta, D. Sing, I. Snellen, F. Snik, J. Spronck, E. Stempels, X. Sun, S. S. Tschudi, and J. Young, in *Ground-based and Airborne Instrumentation for Astronomy VI*, Vol. 9908, edited by C. J. Evans, L. Simard, and H. Takami, *International Society for Optics and Photonics (SPIE)*, (2016) pp. 1949 – 1961.
- [97] Szentgyorgyi, A., J. Bean, B. Bigelow, A. Bouchez, M.-Y. Chun, J. D. Crane, H. Epps, I. Evans, J. Evans, A. Frebel, G. Furesz, A. Glenday, D. Guzman, , T. Haref, J.-G. Jang, B.-H. Jang, U. Jeong, A. Jordan, C.-H. Li, M. Lopez-Morales, K.-M. Kim, J. Kim, K. M. Cracken, B. McLeod, M. Mueller, J. Nah, T. Norton, H. Oh, J. S. Ohd, M. Ordway, B.-G. Park, C. Park, S.-J. Park, D. Phillips, D. Plummer, W. Podgorski, F. Rodler, A. Seifahrt, K.-M. Tak, A. Uomotof, M. A. V. Dam, R. Walsworth, Y. S. Yu, and I.-S. Yuk, *SPIE Proc* **9147-78** (2014).
- [98] D. S. Hayes and D. Latham, *The Astrophysical Journal* **197**, 593 (1975).
- [99] C. W. Allen, *asqu* (1973).

- [100] G. V. Rozenberg, *Twilight* (Springer, 1966).
- [101] D. Deming, F. Espenak, D. E. Jennings, J. W. Brault, and J. Wagner, **316**, 771 (1987).
- [102] G. R. Davies, W. J. Chaplin, Y. Elsworth, and S. J. Hale, *Monthly Notices of the Royal Astronomical Society* **441**, 3009 (2014), <https://academic.oup.com/mnras/article-pdf/441/4/3009/4055353/stu803.pdf> .
- [103] A. T. Young and W. M. Irvine, **72**, 945 (1967).
- [104] D. F. Gray, *The observation and analysis of stellar photospheres* (Cambridge University Press, 2005).
- [105] Reiners, A., Lemke, U., Bauer, F., Beeck, B., and Huke, P., *A&A* **595**, A26 (2016).
- [106] H. B. Snodgrass and R. K. Ulrich, **351**, 309 (1990).
- [107] M. Rostami, F. Koushanfar, and R. Karri, *Proceedings of the IEEE* **102**, 1283 (2014).
- [108] M. Tehranipoor, H. Salmani, X. Zhang, M. Wang, R. Karri, J. Rajendran, and K. Rosenfeld, *Computer* **44**, 66 (2010).
- [109] Semiconductor Industry Association, *SIA Whitepaper*, Washington DC (2013).
- [110] P. Hoeper and J. Manfredelli, *Defense Science Board Washington DC USA* (2017).
- [111] U. Guin, K. Huang, D. DiMase, J. M. Carulli, M. Tehranipoor, and Y. Makris, *Proceedings of the IEEE* **102**, 1207 (2014).
- [112] J. Balasch, B. Gierlich, and I. Verbauwhede, *2015 IEEE International Symposium on Electromagnetic Compatibility (EMC)* , 246 (2015).
- [113] J. Gaudestad and A. Orozco, *Microelectronics Reliability* **54**, 2093 (2014).
- [114] Y. Tagro, J. J. Yan, D. F. Kimball, H. Ghajari, and D. F. Sievenpiper, in *GOMACTech* (2017).
- [115] A. Orozco, J. Gaudestad, N. Gagliolo, C. Rowlett, E. Wong, A. Jeffers, B. Cheng, F. C. Wellstood, A. B. Cawthorne, and F. Infante (2013).
- [116] A. Orozco, *Magnetic Field Imaging for Electrical Fault Isolation* (2019).
- [117] A. N. Campbell, E. Cole, B. A. Dodd, and R. E. Anderson, *31st Annual Proceedings Reliability Physics 1993* , 168 (1993).
- [118] H. H. Huston and C. P. Clarke, *30th Annual Proceedings Reliability Physics 1992* , 268 (1992).
- [119] Semiconductor Industry Association, (2017).
- [120] B. J. Roth, N. G. Sepulveda, and J. P. Wikswo, *Journal of Applied Physics* **65**, 361 (1989).
- [121] E. V. Levine, M. J. Turner, P. Kehayias, C. A. Hart, N. Langellier, R. Trubko, D. R. Glenn,

- R. R. Fu, and R. L. Walsworth, *Nanophotonics* **8**, 1945 (2019).
- [122] D. R. Glenn, R. R. Fu, P. Kehayias, D. Le Sage, E. A. Lima, B. P. Weiss, and R. L. Walsworth, *Geochemistry, Geophysics, Geosystems* **18**, 3254 (2017).
- [123] J. Taylor, P. Cappellaro, L. Childress, L. Jiang, D. Budker, P. Hemmer, A. Yacoby, R. Walsworth, and M. Lukin, *Nature Physics* **4**, 810 (2008).
- [124] A. Gruber, A. Dräbenstedt, C. Tietz, L. Fleury, J. Wrachtrup, and C. v. Borczyskowski, *Science* **276**, 2012 (1997).
- [125] J. R. Maze, P. L. Stanwix, J. S. Hodges, S. Hong, J. M. Taylor, P. Cappellaro, L. Jiang, M. V. G. Dutt, E. Togan, A. S. Zibrov, A. Yacoby, R. L. Walsworth, and M. D. Lukin, *Nature* **455**, 644 (2008).
- [126] G. Balasubramanian, I. Y. Chan, R. Kolesov, M. Al-Hmoud, J. Tisler, C. Shin, C. Kim, A. Wojcik, P. R. Hemmer, A. Krueger, T. Hanke, A. Leitenstorfer, R. Bratschitsch, F. Jelezko, and J. Wrachtrup, *Nature* **455**, 648 (2008).
- [127] L. M. Pham, D. Le Sage, P. L. Stanwix, T. K. Yeung, D. Glenn, A. Trifonov, P. Cappellaro, P. R. Hemmer, M. D. Lukin, H. Park, A. Yacoby, and R. L. Walsworth, *New Journal of Physics* **13**, 045021 (2011).
- [128] M. J. H. Ku, T. X. Zhou, Q. Li, Y. J. Shin, J. K. Shi, C. Burch, H. Zhang, F. Casola, T. Taniguchi, K. Watanabe, P. Kim, A. Yacoby, and R. L. Walsworth, (2019), [arXiv:1905.10791](https://arxiv.org/abs/1905.10791) .
- [129] J. F. Barry, M. J. Turner, J. M. Schloss, D. R. Glenn, Y. Song, M. D. Lukin, H. Park, and R. L. Walsworth, *Proceedings of the National Academy of Sciences* **113**, 14133 (2016).
- [130] A. Nowodzinski, M. Chipaux, L. Toraille, V. Jacques, J.-F. Roch, and T. Debuisschert, *Microelectronics Reliability* **55**, 1549 (2015).
- [131] D. A. Simpson, J.-P. Tetienne, J. McCoey, K. Ganesan, L. T. Hall, S. Petrou, R. E. Scholten, and L. C. L. Hollenberg, *Scientific Reports* **6**, 22797 (2016).
- [132] D. Le Sage, K. Arai, D. R. Glenn, S. J. DeVience, L. M. Pham, L. Rahn-Lee, M. D. Lukin, A. Yacoby, A. Komeili, and R. L. Walsworth, *Nature* **496**, 486 (2013).
- [133] R. R. Fu, B. P. Weiss, E. A. Lima, R. J. Harrison, X.-N. Bai, S. J. Desch, D. S. Ebel, C. Suavet, H. Wang, D. Glenn, D. Le Sage, T. Kasama, R. L. Walsworth, and A. T. Kuan, *Science* **346**, 1089 (2014).
- [134] E. De Mulder, *PhD Thesis* (2010).
- [135] L. Sauvage, S. Guilley, and Y. Mathieu, *ACM Transactions on Reconfigurable Technology and Systems (TRETs)* **2**, 4 (2009).
- [136] A. Horsley, G.-X. Du, and P. Treutlein, *New Journal of Physics* **17**, 112002 (2015).
- [137] J. M. Schloss, J. F. Barry, M. J. Turner, and R. L. Walsworth, *Phys. Rev. Applied* **10**, 034044

- (2018).
- [138] M. W. Doherty, N. B. Manson, P. Delaney, F. Jelezko, J. Wrachtrup, and L. C. L. Hollenberg, *Physics Reports* **528**, 1 (2013).
 - [139] P. Kehayias, M. J. Turner, R. Trubko, J. M. Schloss, C. A. Hart, M. Wesson, D. R. Glenn, and R. L. Walsworth, *Phys. Rev. B* **100**, 174103 (2019).
 - [140] See Supplemental Material at [URL will be inserted by publisher] for additional information on further measurements, integrated circuit information, magnetometer performance, and machine learning details.
 - [141] J. F. Barry, J. M. Schloss, E. Bauch, M. J. Turner, C. A. Hart, L. M. Pham, and R. L. Walsworth, *Rev. Mod. Phys.* **92**, 015004 (2020).
 - [142] F. Dolde, H. Fedder, M. W. Doherty, T. Nöbauer, F. Rempp, G. Balasubramanian, T. Wolf, F. Reinhard, L. C. L. Hollenberg, F. Jelezko, and J. Wrachtrup, (2011), [10.1038/NPHYS1969](https://arxiv.org/abs/10.1038/NPHYS1969).
 - [143] E. Bauch, C. A. Hart, J. M. Schloss, M. J. Turner, J. F. Barry, P. Kehayias, S. Singh, and R. L. Walsworth, *Phys. Rev. X* **8**, 031025 (2018).
 - [144] X. Zhang and M. Tehranipoor, *2011 Design, Automation & Test in Europe*, 1 (2011).
 - [145] M. Mandal and B. C. Sarkar, *Indian Journal of Pure and Applied Physics* **48**, 136 (2010).
 - [146] V. M. Acosta, E. Bauch, M. P. Ledbetter, A. Waxman, L.-S. Bouchard, and D. Budker, *Phys. Rev. Lett.* **104**, 070801 (2010).
 - [147] V. V. Dobrovitski, A. E. Feiguin, D. D. Awschalom, and R. Hanson, *Phys. Rev. B* **77**, 245212 (2008).
 - [148] L. Shao, R. Liu, M. Zhang, A. V. Shneidman, X. Audier, M. Markham, H. Dhillon, D. J. Twitchen, Y.-F. Xiao, and M. Lončar, *Advanced Optical Materials* **4**, 1075 (2016).
 - [149] A. Przybylski, B. Thiel, J. Keller-Findeisen, B. Stock, and M. Bates, *Scientific Reports* **7** (2017).
 - [150] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, *Journal of Machine Learning Research* **12**, 2825 (2011).
 - [151] K. P. F.R.S., *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **2**, 559 (1901).
 - [152] H. Hotelling, *Biometrika* **28**, 321 (1936).
 - [153] C. Cortes and V. Vapnik, *Machine Learning* **20**, 273 (1995).
 - [154] P. Kocher, J. Jaffe, and B. Jun, in *Advances in Cryptology — CRYPTO' 99*, edited by M. Wiener (Springer Berlin Heidelberg, Berlin, Heidelberg, 1999) pp. 388–397.

- [155] S. B. Örs, E. Oswald, and B. Preneel, in *Cryptographic Hardware and Embedded Systems*, edited by C. D. Walter, Ç. K. Koç, and C. Paar (Springer Berlin Heidelberg, Berlin, Heidelberg, 2003) pp. 35–50.
- [156] U. Rührmair, X. Xu, J. Sölter, A. Mahmoud, M. Majzoubi, F. Koushanfar, and W. Burleson, in *Cryptographic Hardware and Embedded Systems*, edited by L. Batina and M. Robshaw (Springer Berlin Heidelberg, Berlin, Heidelberg, 2014) pp. 476–492.
- [157] A. Ravi, N. Langellier, D. F. Phillips, M. Buschmann, B. R. Safdi, and R. L. Walsworth, *Phys. Rev. Lett.* **123**, 091101 (2019).
- [158] G. Bertone and D. Hooper, *Rev. Mod. Phys.* **90**, 045002 (2018).
- [159] J. I. Read, *J. Phys. G: Nucl. Part. Phys.* **41**, 063101 (2014).
- [160] N. Meunier, A.-M. Lagrange, and M. Desort, *Astron. Astrophys.* **519**, A66 (2010).
- [161] S. Roberts, M. Osborne, M. Ebden, S. Reece, N. Gibson, and S. Aigrain, *Philos. Trans. Royal Soc. A* **371**, 20110550 (2013).
- [162] J. Nocedal and S. J. Wright, *Numerical Optimization*, edited by P. Glynn and S. M. Robinson (Springer, 1999) pp. 136–138.