

Georgia State University

ScholarWorks @ Georgia State University

Physics and Astronomy Dissertations

Department of Physics and Astronomy

Fall 12-12-2022

Expanding the Search for the Pulse of Jupiter

Deborah Gullledge

Follow this and additional works at: https://scholarworks.gsu.edu/phy_astr_diss

Recommended Citation

Gullledge, Deborah, "Expanding the Search for the Pulse of Jupiter." Dissertation, Georgia State University, 2022.

doi: <https://doi.org/10.57709/32559304>

This Dissertation is brought to you for free and open access by the Department of Physics and Astronomy at ScholarWorks @ Georgia State University. It has been accepted for inclusion in Physics and Astronomy Dissertations by an authorized administrator of ScholarWorks @ Georgia State University. For more information, please contact scholarworks@gsu.edu.

Expanding the Search for the Pulse of Jupiter

by

Deborah J. Gulledge

Under the Direction of Stuart Jefferies, Ph.D.

A Dissertation Submitted in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

in the College of Arts and Sciences

Georgia State University

2022

ABSTRACT

Although we have learned much about Jupiter throughout history, many questions still remain, with one in particular that we seek to answer: does a solid core lie at the center of the planet, or is it gaseous and metallic throughout? To uncover this mystery, we turn to seismology, a sole remaining way to directly and unambiguously probe the deepest interior of Jupiter. We have developed a novel set of instrumentation, called **PMODE**—the **P**lanetary **M**ultilevel **O**scillations & **D**ynamics **E**xperiment, a triple-channel instrument to probe the Jovian interior through seismic techniques. Here, we focus on the 770 nm potassium Doppler imager. This channel is developed to measure the Doppler shift of reflected sunlight off the Jovian clouds; these shifts contain information on the densities deep within the planet. With PMODE, we have collected 24 nights of observations in August 2020 from the 3.6m AEOS telescope, located on Mount Haleakalā, HI.

From these observations, we have produced a sensitive measurement of the high-frequency components of the Jovian zonal winds, alleviating contradiction between the sole previous Doppler measurement and the well studied cloud tracking measurement. We provide evidence that the PMODE Doppler measurement, combined with a separate measurement sensitive to the low-spatial frequency components (such as the previous Doppler measurement), will accurately reproduce the profile from feature tracking models. We refute the previous report of detection of Jovian global modes with amplitudes of 50 cm s^{-1} , and instead place a constraint of 4.5 cm s^{-1} for amplitudes in the mode of historical interest. This new constraint eliminates rock storms as an excitation mechanism present in the Jovian interior, implying that any existing oscillations will likely have amplitudes below 1 cm s^{-1} . Finally, we discuss avenues to detect these smaller amplitudes in the future, and conclude that inclusion of a Doppler velocimeter on a Jovian orbiter may provide the only chance at an unambiguous detection of the global modes of Jupiter.

INDEX WORDS: Jupiter, Planetary Seismology, Planetary Interiors, Planetary Atmospheres, Astronomical Instrumentation

Copyright by
Deborah J. Gullledge
2022

Expanding the Search for the Pulse of Jupiter

by

Deborah J. Gulledge

Committee Chair:

Stuart Jefferies

Committee:

Fabien Baron

Todd Henry

Neil Murphy

Electronic Version Approved:

Office of Graduate Studies

College of Arts and Sciences

Georgia State University

December 2022

DEDICATION

To my mother, Kathy, and my sisters, Laura and Whitney. I love you endlessly and eternally.

In loving memory of my father, Ellis Preston Gulledge, Jr.—I am proud to know that you live on in my sisters and I.

ACKNOWLEDGMENTS

My advisor, Stuart. Thank you for taking a chance on a wide-eyed grad student who waltzed into your office declaring her biggest dreams, and for giving me every opportunity to make those dreams a reality. You have pushed me to become a much better scientist and I am immensely thankful for it.

My committee members: Todd, thank you for bringing me to GSU in the first place, for sharing a love of planets with me, and providing not only the academic support, but also the emotional support I needed to thrive; Fabien, for not only acting as a co-advisor but for telling the best jokes I've ever heard; and Neil, for the long hours and late nights staying up to help me observe, and the kindness you provided along the way.

My Maui mentors: Ryan, for the unconditional welcome into an amazing team working on the craziest project I've ever been a part of, for the patience in bringing me up to speed, for the opportunities you've given me, and the sushi and laughs along the way. Chris, for the countless hours we spent working on LANDIT together, for the career and life advice, and for not only being a friend and mentor, but a father figure for me along the way. And Cody—this dissertation exists because of you, in particular. You've kept me curious on this journey and given me endless patience, guidance, and support throughout it. I couldn't thank you enough if I had all the time granted by the Ash Twin Project. For all my mentors—I have learned more from you all than I could ever put into words, and I am tremendously grateful for the knowledge you have shared with me.

My teachers throughout the years—in science, Dr. Smith, Dr. Mayo, Dr. Longhurst, and Ms. Livengood, who fueled my curiosity for the world around me, and in literature, Dr. Cervelli, Mr. Welch, and Mr. Yuhas, for reminding me to search for the poetry in the world, for teaching me that science and art are inexorably linked.

My incredible friends, of which there are more I am thankful to know than I could ever list. An extra thank-you to my longest friends: Joy, for loving me since we shared a crib together and every day that followed; Clarissa, for being the greatest roommate through undergrad, for our cross-country road trips and impromptu musicals; AJ, for your constant friendship from high school marching band together through growing into the person I am today; and Jenn, for being the other mama to my babies and for showing me the beauty that comes from living a life that makes *yourself* the happiest. My Doctor- and proto-Doctor friends from GSU: Katherine, for being my first and best friend in Atlanta and my twin in the program; Rachael, for always answering my panicked phone calls, giving the best advice, and making me feel so loved; Justin, for being the most encouraging person I know and for the Anatolia trips when things got too stressful; Mitch, for being not only one of the most amazing friends, but also for the mentorship through all these years—I've looked up to you since my first day in the program; and Emma, for being my perfect Astropal and twin soul in the department. Drew and Olivia, for the beautiful friendships that grew from a spontaneous trip to Florida, and for reminding me to find peace underwater throughout this journey. Finally, my Destiny clan—thank you for the raids, the laughs, and making me take the all-important breaks. Getting Rivensbane with y'all is one of my fondest memories.

And my family. My cats, Baelfire and Arkham, who walked their little pawprints all over my heart with impressions that will last a lifetime and who have given me infinite comfort and unconditional love over my academic career.

My sisters, the other pieces to my heart and the best friends I could dream of. To my baby sister, Laura Beth. You are an embodiment of all that is good, all that is kind, all that is beautiful in this world. I love you, tiny bear. Thank you for being the first person I turn to when I need someone to cry to, for sharing your optimism and love for life on my journey through school. I am the scientist I am today because you always told me you looked up to me—making you proud has been the most beautiful motivation I could ever ask for.

And Whit. You may be my little sister, too, but I have looked up to *you* for as long as I can remember. You are strength and sunshine all rolled into one. Thank you for being an unwavering pillar of support for this family, for giving me the exact encouragement I always need, and making me laugh when things feel the darkest. I am glad we've been able to fight through school together, and I can't wait to meet Whitney Gulledge, M.D.—you amaze me.

Finally, my Mama—you have shown me what the words strength and love mean. Everything I do, and everything I am, is because of you. You have given me the freedom to fly, and the safety net to catch me when I fall. You have taught me to find the beauty in every situation, and how to handle hardships with grace. Your unending warmth and support and your encouragement of my curiosity of the world around me are the reason I am writing this dissertation today. Thank you for being not only my mother, but also my greatest friend. I love you.

*“I hope you still feel small when you stand beside the ocean,
Whenever one door closes I hope one more opens,
Give the heavens above more than just a passing glance,
and when you get the choice to sit it out, or dance—
I hope you dance.”*

-Lee Ann Womack, 2000

Deborah Gullledge gratefully acknowledges support from the National Science Foundation through the Graduate Research Fellowship Program (GRFP). This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 1937956. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

TABLE OF CONTENTS

ACKNOWLEDGMENTS	v
LIST OF TABLES	xiv
LIST OF FIGURES	xv
LIST OF ABBREVIATIONS	xviii
1 INTRODUCTION	1
1.1 On Planetary Formation—The Grand Question	1
1.2 Space-Based Attempts to Probe the Jovian Interior	3
1.2.1 <i>Current Understanding of the Interior Structure of Jupiter</i>	5
1.2.2 <i>An Aside—The Saturnian Interior</i>	7
1.3 Implications on Planetary Formation from Gas Giant Interiors	8
1.4 DiOSEISMology: Probing the Deep Jovian Interior with Seismology	10
1.4.1 <i>Inheritance from Helioseismology</i>	13
1.4.2 <i>Historical DiOSEISMology Efforts</i>	15
1.5 Jovian Atmospheric Science	19
1.5.1 <i>The Jovian Zonal Wind Profile</i>	19
1.5.2 <i>The Odd Gravitational Moments as a Probe of Atmospheric Dynamics</i>	22
1.5.3 <i>Applications of Doppler Velocimetry to Non-Jovian Atmospheric Science</i>	24
1.6 Remaining Questions	25
1.6.1 <i>Remaining Questions Regarding Jovian Atmospheric Dynamics</i>	25
1.6.2 <i>Remaining Questions Regarding Jovian Interior Structure</i>	27
1.7 Dissertation Overview	28

2	ON DEVELOPING INSTRUMENTATION FOR MEASUREMENT OF JOVIAN OSCILLATIONS	30
2.1	Instrumental Concept	30
2.2	PMODE: the Doppler Imager Channels	33
2.2.1	<i>MOF Optical Components and Transmission Path</i>	35
2.2.2	<i>Doppler Velocity Spatial Signal on the Jovian Disk</i>	37
2.3	Full Experiment Optical Design	41
2.3.1	<i>PMODE Full Optical Design—889 nm Polarimeter Path</i>	45
2.3.2	<i>PMODE Full Optical Design—589 nm Sodium Path</i>	46
2.3.3	<i>PMODE Full Optical Design—770 nm Potassium Path</i>	48
2.3.4	<i>PMODE Full Optical Design—Deviations from the Original Design</i>	49
2.4	Sodium Channel Detector Artefacts	51
2.5	Theoretical Performance	54
3	ON AN OBSERVATIONAL CAMPAIGN FROM MT. HALEAKALĀ, HI	60
3.0.1	<i>Observing Conditions</i>	61
3.0.2	<i>On-Sky Optimization</i>	64
3.1	PMODE Data Reduction	67
3.1.1	<i>Standard Calibrations</i>	69
3.1.2	<i>Moon Flats</i>	73
3.1.3	<i>Field De-Rotation and Image Registration</i>	74
3.2	PMODE Residual Generation	78
3.2.1	<i>Additional Velocity Calibrations</i>	78
3.2.2	<i>Dopplergram Cleaning</i>	83
3.3	Calibration of Raw MOF Data to Velocity Units	86
3.3.1	<i>Original Calibration of Raw MOF Data to Velocity Units</i>	86
3.3.2	<i>Improved Calibration of Raw MOF Data to Velocity Units</i>	88
3.4	Spherical Harmonic Multiplication	90

3.5	Lomb-Scargle Periodogram Calculation to Obtain Final Amplitude Spectra	92
3.6	Analysis of Spherical Harmonics	98
3.6.1	<i>Spherical Harmonic Leakage</i>	98
3.6.2	<i>Determination of Final Set of Modes for Analysis</i>	100
3.7	Additional Noise Sources	108
3.7.1	<i>Tracking and Image Registration</i>	108
3.7.2	<i>Distortion Noise</i>	109
3.7.3	<i>Temperature Noise</i>	109
4	ON SIGNIFICANT RESULTS FROM THE PMODE CAMPAIGN . . .	111
4.1	Detection of Earth's Rotational Signal within the PMODE Data . .	111
4.2	A New Measurement of the Jovian Zonal Winds	114
4.2.1	<i>Extracting the Zonal Wind Profile from PMODE Data</i> . . .	115
4.2.2	<i>Final PMODE Measurement of the Zonal Wind Profile</i> . .	119
4.2.3	<i>Comparison of PMODE ZWP with Historical Measurements</i>	122
4.2.4	<i>New Implications for Jovian Wind Velocities from the PMODE Measurement</i>	129
4.3	Characterization of Power Spectrum Noise Levels	132
4.3.1	<i>"Random-Permutation" (RP) Noise Characterization</i> . . .	132
4.3.2	<i>"Background-Fitting" (BF) Noise Characterization</i>	133
4.4	Final PMODE Amplitude Spectra: The (1,0) Mode	141
4.5	Final PMODE Amplitude Spectra for All Defined Modes of Interest	143
4.5.1	<i>Final Amplitude Spectra: Low-Order Modes</i>	143
4.5.2	<i>Final Amplitude Spectra: $l = m$ Modes</i>	148
4.5.3	<i>Final Amplitude Spectra: Hemispherically Antisymmetric Modes</i>	151
4.5.4	<i>Review of Presented Amplitude Spectra and Implications for Remaining Modes</i>	154

5	ON PRESENT COMPLICATIONS AND FUTURE PATHS FOR DETECTING THE JOVIAN OSCILLATIONS	155
5.1	Further Validation of PMODE Velocimetric Sensitivity with an Astronomical Scene Generator	156
5.1.1	<i>Computer-Generated Holography with a Spatial Light Modulator</i>	158
5.1.2	<i>First Attempt for Determination of Input Patterns for the SLM</i>	160
5.1.3	<i>Optical Path for the Astronomical Scene Generator</i>	162
5.1.4	<i>Initial Complications and Second Attempt for Determination of Input Patterns for the SLM</i>	165
5.1.5	<i>Third Attempt for Determination of Input Patterns for the SLM</i>	170
5.2	Future Instrumentation Improvements	171
5.2.1	<i>Increasing Spatial Coverage on the Jovian Disk</i>	172
5.2.2	<i>A Novel Imaging Shack-Hartmann Wavefront Sensor</i>	173
5.3	Future Paths for Detecting the Jovian Global Modes of Oscillation	175
5.3.1	<i>Adaptation of the PMODE Instrumentation for Deployment to the Geographic South Pole</i>	177
5.3.2	<i>LANDIT Thermal Testing</i>	181
5.3.3	<i>Complications for a 0.5m Polar Telescope</i>	183
5.3.4	<i>A Multi-Site Polar Network for Detecting Jovian Oscillations</i>	186
5.3.5	<i>The Final Frontier: Detecting Jovian Oscillations from a Dedicated Space-Based Doppler Imager</i>	187
6	CONCLUSIONS	190
	Appendices	194
A	MATLAB Code to Reduce PMODE Data	195
B	MATLAB Code to Generate PMODE Residuals	206
C	MATLAB Code to Multiply Residuals by Spherical Harmonics	212
D	MATLAB Code to Calculate Final Amplitude Spectra	216

REFERENCES	222
----------------------	-----

LIST OF TABLES

Table 2.1	Optical Transmission through a PMODE MOF Arm	57
Table 3.1	Summer 2020 PMODE Observational Campaign Statistics	64
Table 4.1	Average Noise Values for $Y_l^m = (0, 0)$ to $Y_l^m = (5, 5)$ Power Spectra from the PMODE Campaign	137
Table 4.2	Average Noise Values for $Y_l^m = (6, -6)$ to $Y_l^m = (8, -1)$ Power Spectra from the PMODE Campaign	138
Table 4.3	Average Noise Values for $Y_l^m = (8, 0)$ to $Y_l^m = (10, -3)$ Power Spectra from the PMODE Campaign	139
Table 4.4	Average Noise Values for $Y_l^m = (10, -2)$ to $Y_l^m = (11, 11)$ Power Spectra from the PMODE Campaign	140

LIST OF FIGURES

Figure 2.1	Spectrum of Jupiter with PMODE Channels Highlighted	33
Figure 2.2	Transmission of Light Through a Magneto-Optical Filter	34
Figure 2.3	PMODE: Collected Observations to Dopplergrams	37
Figure 2.4	Visualization of the 770 nm K Absorption Line as a Function of both Wavelength and Doppler Shift	39
Figure 2.5	Visualization of the 889 nm Na DI Absorption Line as a Function of both Wavelength and Doppler Shift	41
Figure 2.6	Simulated K & Na MOF Sensitivity for PMODE	42
Figure 2.7	The Small-Footprint Optical Breadboard for the PMODE Instrument	43
Figure 2.8	Diagram of the Optical Path for the PMODE Instrument	44
Figure 2.9	Optical Diagram of the Image Plane vs. the Pupil Plane	45
Figure 2.10	Example of the Etaloning Artefacts in the Na Dataset	53
Figure 3.1	Temperature Determination Process for the K and Na MOFs	66
Figure 3.2	Final Determined Temperatures for the K and Na MOFs	68
Figure 3.3	Basic PMODE Data Reduction Frames and Steps	70
Figure 3.4	Change in Absorption Line Center from Earth-Jupiter Relative Velocity	80
Figure 3.5	PMODE: Dopplergrams to Residual Images	83
Figure 3.6	PMODE Example Residual Cutoff	84
Figure 3.7	PMODE: Dopplergram Absorption Line Cleaning Process	85
Figure 3.8	MOF Unit Calculation—Modified Technique for Converting Data Num- bers to Velocity	89
Figure 3.9	Example Set of Projected Spherical Harmonics	91
Figure 3.10	Tukey Window Example	93

Figure 3.11 Residual Smoothing and Tukey Windowing	95
Figure 3.12 Full Campaign Residual Outlier Rejection	96
Figure 3.13 Spherical Harmonic Leakage Matrices for the PMODE Observational Campaign	101
Figure 3.14 Comparison of the $Y_\ell^m = (11, -11)$ Mode to the Width of the PMODE 770 nm K Absorption Line	105
Figure 3.15 Spherical Harmonic Leakage Matrix— $Y_\ell^m = (0, 0)$ to $Y_\ell^m = (11, \pm 11)$	106
Figure 4.1 Earth’s Velocimetric Rotation Signal from the PMODE Data	113
Figure 4.2 Earth’s Velocimetric Rotation Signal from the PMODE Data	114
Figure 4.3 Filtering Procedures to Reveal the Jovian Zonal Winds	118
Figure 4.4 PMODE Velocimetric Measurement of the Jovian Zonal Wind Profile	119
Figure 4.5 PMODE Zonal Wind Profile Measurement Overlaid on the Jovian Disk	121
Figure 4.6 PMODE Doppler Zonal Wind Profile Compared to Previous Doppler Zonal Wind Profile	123
Figure 4.7 PMODE Doppler Zonal Wind Profile Compared to <i>Hubble</i> Cloud Tracking Zonal Wind Profile	126
Figure 4.8 HST Zonal Wind Profile Decomposed Into Legendre Polynomials	128
Figure 4.9 PMODE Zonal Wind Profile Added to Low-Order Legendre Polynomial HST Decomposition	131
Figure 4.10 “Random-Permutation” (RP) Background Noise Example	134
Figure 4.11 “Background-Fitting” (BF) Noise Example	135
Figure 4.12 PMODE (1,0) Power Spectrum for Comparison with Previous Results	142
Figure 4.13 PMODE Amplitude Spectra for the Low-Order $Y_\ell^m = (0, 0)$ to $Y_\ell^m = (1, 1)$ Modes	144
Figure 4.14 PMODE Amplitude Spectra for the Low-Order $Y_\ell^m = (2, -2)$ to $Y_\ell^m = (2, 2)$ Modes	145
Figure 4.15 PMODE Amplitude Spectra for the Low-Order $Y_\ell^m = (3, -3)$ to $Y_\ell^m = (3, 0)$ Modes	146

Figure 4.16 PMODE Amplitude Spectra for the Low-Order $Y_\ell^m = (3, 1)$ to $Y_\ell^m = (3, 3)$ Modes	147
Figure 4.17 PMODE Amplitude Spectra for the $\ell = m = \pm 4, 5, 6$ Modes	149
Figure 4.18 PMODE Amplitude Spectra for the $\ell = m = \pm 7, 8, 9$ Modes	150
Figure 4.19 PMODE Amplitude Spectra for the $\ell = m = \pm 10, 11$ Modes	151
Figure 4.20 PMODE Amplitude Spectra for the Hemispherically Antisymmetric Modes, $Y_\ell^m = (4, \pm 1), Y_\ell^m = (4, \pm 3)$ Modes	152
Figure 4.21 PMODE Amplitude Spectra for the Hemispherically Antisymmetric Modes, $Y_\ell^m = (5, 0), Y_\ell^m = (5, \pm 2), Y_\ell^m = (5, \pm 4)$ Modes	153
Figure 5.1 Liquid Crystal Manipulation for Phase Modulation with an SLM	159
Figure 5.2 Representation of the Zeroth-Order Spot for an SLM	160
Figure 5.3 First Optical Path for the SLM Object Simulator	164
Figure 5.4 Output Images from the First Optical Path for the SLM Object Generator	165
Figure 5.5 Simulated Outputs for Phase Calculation for the SLM Object Generator	169
Figure 5.6 Phase-Only HST Image Simulation from the SLM Object Generator .	171
Figure 5.7 Phase-Only Jupiter Image Simulation from the SLM Object Generator	172
Figure 5.8 The Assembled LANDIT Telescope	179
Figure 5.9 LANDIT Thermal Housing Deep Freezer Test	183

LIST OF ABBREVIATIONS

AEOS	Advanced Electro Optical System
AI/ML	Artificial Intelligence / Machine Learning
AO	Adaptive Optics
AU	Astronomical Unit
BF	Background-Fitting
CAD	Computer Aided Design
CCD	Charge-Coupled Device
CGH	Computer-Generated Holography
DN	Data Number
EMCCD	Electron Multiplying Charge-Coupled Device
ESA	European Space Agency
EZ	Equatorial Zone
FFT	Fast Fourier Transform
FLI	Finger Lakes Instrumentation
FWHM	Full Width at Half Maximum
GRS	Great Red Spot
HST	Hubble Space Telescope
ISHWFS	Imaging Shack-Hartmann Wavefront Sensor
JPL	Jet Propulsion Laboratory

JIVE	Jovian Interiors from Velocimetry Experiment
JOVIAL	Jovian Oscillations through Velocity Images At several Longitudes
LANDIT	Long-duration Antarctic Night and Day Imaging Telescope
MAD	Median Absolute Deviations
MJD	Modified Julian Date
MOF	Magneto-Optical Filter
NASA	National Aeronautics and Space Administration
NEB	Northern Equatorial Belt
NIST	National Institute of Standards and Technology
NOAA	National Oceanic and Atmospheric Administration
NTB	Northern Temperate Belt
NTrZ	Northern Tropical Zone
OE	Optical Element
OPAL	Outer Planet Atmospheres Legacy
PEARL	Polar Environment Atmospheric Research Laboratory
PID	Proportional Integral Derivative
PMODE	Planetary Multilevel Oscillations and Dynamics Experiment
PSF	Point-Spread Function
RP	Random-Permutation
SDSS	Sloan Digital Sky Survey
SHWFS	Shack-Hartmann Wavefront Sensor

SLM	Spatial Light Modulator
S/N	Signal-to-Noise
SNR	Signal-to-Noise Ratio
SYMPA	Seismographic Imaging Interferometer for Monitoring of Planetary At-
	mospheres
ZWP	Zonal Wind Profile

CHAPTER 1

INTRODUCTION

1.1 On Planetary Formation—The Grand Question

One of the most burning and long-standing questions that remains in the astronomical community today is also one of the most fundamental—how did our Solar System, itself, form? Did the planets grow from the smallest, micrometer-sized dust particles slowly crashing into one another and combining, this process repeating until the small mass assumes enough gravity to accrete more surrounding material from the protoplanetary disk, known as the core accretion process? (Bodenheimer & Pollack (1986), Pollack et al. (1996), Alibert et al. (2005)) And if this core accretion process occurs, what are the average sizes of the accreted material—were they smaller, millimeter-to-centimeter sized particles (pebble accretion) that accrete on much shorter timescales than the alternative option: larger-than-centimeter to asteroid-sized solids (planetesimal accretion), or some combination of the two? (Johansen & Lambrechts 2017).

Or instead, did the large, gaseous protoplanetary disk, rife with density non-uniformities, collapse down from these overdensities into planetary bodies in the disk instability manner, similarly to stellar formation? (Boss (1997), Mayer et al. (2002)) Perhaps the most intuitive way to discern which of these theories is correct (accepting the inability to directly look back in time and observe our own formation process) is by revealing the structure which lies deep within the Gas Giants that call our Solar System home. Specifically, examining the radial distribution of heavy elements will allow us to determine if a solid core lies deep within the

enshrouding gaseous envelopes. And if this core exists, is it compact, with a sharp, well-defined boundary, or is it diffuse and diluted, with an extended, mixed gradient between the pure, heavy element core and the gaseous outer envelope?

Direct evidence revealing the existence of a solid core (regardless of whether this core displays a defined boundary, or a diffuse one) holds the potential to entirely refute the disk instability formation theory. When considering an originally homogeneous distribution of elements (such as a protoplanet which may arise from the disk instability formation theory) with a Jovian mass and radius, and assuming typical diffusive settling driven by differences in molecular weights, the necessary duration for this settling to create a heavy-element core requires a timescale longer than the age of the Universe. If the additional impact of re-mixing by convection while this settling is occurring is included in this calculation, this timescale for settling into a core becomes even longer (Stevenson 2020). Therefore, core accretion is the current most-agreed-upon theory, and is comprised of three primary iterations (Helled et al. 2022):

- Phase 1: Accretion of primarily, pure solid heavy elements *via* the pebble accretion mechanism, with a low rate of light element (H and He) accretion.
- Phase 2: Heavy element accretion continues, although larger planetesimals comprise the bulk of the accreted matter, at a fairly rapid rate to prevent planetary cooling and subsequent runaway gas accretion. Here, the slow accretion of light elements also begins. The combination of light and heavy elements accreting together causes compositional gradients to begin to form during this stage—once the core regime is

on the order of a few Earth masses, subsequently accreted solids oft dissolve in the envelope before sinking to the core.

- Phase 3: The light element mass exceeds the heavy element mass, leading to runaway gas accretion. Theoretically, this occurs at a Jovian mass of ~ 100 Earth masses. The majority of the mass gained by the planet occurs within this stage.

1.2 Space-Based Attempts to Probe the Jovian Interior

To answer the established burning questions and gain further insight into planetary formation processes, multiple spacecraft have been launched to orbit both Jupiter and Saturn—most profoundly, the *Juno* mission for Jupiter (Bolton et al. 2017), and the *Cassini* mission for Saturn (Edgington & Spilker 2016). The interiors and atmospheres of these Gas Giants have been characterized through the measurement of gravitational moments, or deviations from the spherically symmetric gravitational field which would be produced by a perfectly spherical planet with mass and densities distributed evenly throughout.

However, as neither Jupiter nor Saturn fits this ideal condition—both have variations in their radial distribution of heavy elements—this gravity field is slightly distorted. These distortions are described with terms referred to as J_n -components, which can be thought of as deviations from a sphere in patterns similar to those represented by spherical harmonics. A hydrodynamic, stable Jupiter with no dynamical effects (e.g., tides, oscillations, winds) will have a gravitational potential represented by the following:

$$V(r, \theta) = \frac{GM_J}{r} \left[1 - \sum_{n=1} J_{2n} \left(\frac{R_J}{r} \right)^{2n} P_{2n}(\cos \theta) \right] \quad (1.1)$$

Here, V represents external gravitational potential, dependent on the radius from the center of the planet (r), and the co-latitude (θ). The J_{2n} term are the dimensionless numbers representing the (even) gravitational components—assuming no dynamical effects, these harmonics are the only ones that exist because the only possible perturbations can arise from planetary rotation, which has no variation between the Northern and Southern hemispheres. Finally, P_{2n} represents Legendre polynomials (Stevenson 2020), (Serra et al. 2019). By combining Equation 1.1 with the $V(r, \theta)$ evaluated directly from the Jovian density distribution, we obtain the following representation for the gravitational moments:

$$J_{2n} = \frac{-1}{MR^{2n}} \int \rho(r, \theta) r^{2n} P_{2n}(\cos \theta) d^3r \quad (1.2)$$

where all variables are the same as their counterparts in Equation 1.1. As a result of the r^{2n} term inside the integrand, the higher-order J-terms contain information localized to the outer regions of the planet (Stevenson 2020), a phenomenon displayed in Guillot (2005), Figure 4.

These J-terms will also become smaller as n increases, because as n increases, sensitivity is isolated to higher-order terms, and thus, smaller effects. The complicating factor of these components arises when we consider that de-coupling of these terms from one another remains difficult—a similar measurable effect can be produced by simulating a large change within the deep interior of the planet, or by a small change near the surface—with no information as to which situation is correct. This non-uniqueness and sensitivity to the outer region produces a limitation to the use of the gravitational moments as a probe of the deep interior of the planet, because the density profile can be fit by multiple internal

structure solutions. Varying one model—such as the equation of state, the expected heavy element enrichment, or internal density—can produce similar fits to the *Juno* gravity data with significantly varying interiors and compositions (Helled et al. 2022).

1.2.1 Current Understanding of the Interior Structure of Jupiter

However, these orbital measurements of the gravitational field have allowed for exquisite constraints on the outer regime of the Jovian structure and atmosphere (Bolton et al. 2017), (Folkner et al. 2017), (Iess et al. 2018). The current models providing the most accurate fit to these well-studied outer levels of Jupiter suggest that the internal composition of Jupiter is inhomogeneous, implying that the planet is not fully convective (Helled et al. 2022), that the deep interior is highly enriched with heavy elements, and that the internal structure is stratified, following a structure (moving from the outermost clouds inwards) as such: an outer atmospheric layer consisting of a molecular envelope with a low enrichment of helium and heavy elements, followed by a transition region separating the molecular envelope from the deeper mantle—here, the molecular-metallic transition occurs and helium rain is expected.

Deeper past this transition lies the mantle region, comprised primarily of metallic hydrogen and enriched by some amount of heavy elements and helium; this region gradually becomes more and more enriched by heavy elements to reach a denser, heavy element concentration near the center of the planet (Stevenson 2020). This transition to a denser, heavy “core” region is expected to begin around the inner 50% of the Jovian radius (Wahl et al. (2017), Liu et al. (2019)), comprised of $\sim 30\%$ the total Jovian mass, or around 100 Earth masses (Helled et al. 2022) and predicted to have a “fuzzy” boundary (Helled & Steven-

son 2017) as opposed to a well-defined transition. The majority of models which are most successful at fitting the *Juno* gravity data suggest that this outer “fuzzy core” contains approximately 15 Earth masses of enriched, heavy material, and most interestingly, that deep within this diffuse core lies a solid, pure heavy-element core of 1-5 Earth masses, or $\sim 1\%$ of the Jovian mass, extending to $\sim 10\%$ of the Jovian radius (Helled et al. 2022).

This central pure, heavy-element core is consistent and allowed (although not required, as several models are capable of fitting the gravity data with no small, pure, central core) with current formation models suggesting a core accretion mechanism with a runaway gas accretion period, as discussed in Brouwers & Ormel (2020) and Valletta & Helled (2020). However, the gravity data itself is particularly insensitive to the deepest interior, and as such, any properties of the the pure, heavy element core are poorly constrained (Helled et al. 2022). Valletta & Helled (2020) acknowledge that still, even the origin of the larger, surrounding dilute core present within Jupiter is unknown and must be investigated further.

These current models which provide the best fit to the observed gravity field have no sharp, well-defined boundaries within the planet. Even so, the material is expected to become enriched with heavier elements as we move inwards from the outer molecular envelope of Jupiter towards the core, as defined above. However, this is not believed to result from diffusive settling (as aforementioned, which would take a prohibitively long timescale), but instead, because the ratio of gas-accretion to solid-accretion will increase as the planetary formation epoch progresses (Stevenson 2020), (Helled & Stevenson 2017). These *Juno* data, best fit models, and current results challenge three simple assumptions of Jupiter that were

previously accepted: that it is homogeneously mixed, fully convective, and containing a pure, heavy element core with a clearly defined boundary. That none of these previous assumptions are true suggests that there is still much to learn of the interior of our own Gas Giants.

1.2.2 An Aside—The Saturnian Interior

Although this dissertation focuses primarily on Jupiter, it is interesting to consider Saturn as well when evaluating formation mechanisms and interior structure for giant planets. Saturn’s striking rings provide a bevy of useful information for scientific study. Notably, these rings act as a sensitive seismograph, allowing for the detection of g-modes propagating through Saturn and to its rings (Fuller 2014)—heralding the birth of Kronoseismology, or seismology of Saturn. Here, pressure-modes propagating within the interior of Saturn can create enough gravitational potential to disturb the surrounding rings, creating wave-like spiral patterns within the rings, which can be directly observed. (Unfortunately, Jupiter’s rings are much smaller and dimmer, eliminating this type of seismology as an option for Jovian seismology.) Prior to these detections, it was determined from *Cassini* data that there was evidence of an enhancement of central densities (i.e., a core), within the Saturnian atmosphere (Guillot 2005).

These gravity data from *Cassini*, combined with ring seismology measurements as inputs into sensitive models, produce a precise and sensitive map of the core-envelope transition occurring in the Saturnian interior. They provide also a remarkably tight constraint of Saturn’s density profile, and the mass and radius of the central core: a total heavy-element mass of ~ 19 Earth masses, and an extended radius of ~ 0.32 Saturnian radii (Mankovich &

Fuller 2021). Most interestingly, these findings suggest that Saturn *must* have a diffuse core, as well, with the most likely models hosting a compositional gradient that extends all the way to the central point of the planet.

1.3 Implications on Planetary Formation from Gas Giant Interiors

Evidence suggesting a central core within *both* Jupiter and Saturn is simultaneously fascinating and perplexing. Prior to the discovery of evidence for Saturn’s diffuse core, it was the general assumption that planets likely formed *via* the core-accretion model, with solid, defined boundaries at the interior cores. Jupiter was theorized to be an outlier, with its “fuzzy” core occurring perhaps as the result of a giant impactor, which has the capability to shatter proto-cores within a primordial Jupiter, resulting in the dilute scenario (Liu et al. 2019).

However, the chances of the perfect scenario to produce this outcome are slim—the impactor must be head-on, with a minimum mass of 10 Earth masses, and the occurrence of an impact such as this is very dependent on orbital and situational conditions. Should even one of these conditions change, the impact will no longer produce the diffuse core. If the chances of this occurring once for Jupiter are small, the chances of it occurring for *both* Jupiter and Saturn are microscopic, indeed.

Of course, it must be considered that other effects and dynamical evolution could have occurred during the early period of Solar System formation, as described by the Nice model (Tsiganis et al. 2005). There are two different expansions to the Nice model which may

increase the chances of an impactor striking the Gas Giants to form a “fuzzy core”. The first is known as the “Jumping Jupiter” scenario (Brasser et al. 2009), and suggests that the scattering of an Ice Giant (Uranus or Neptune) inwards results in a rapid movement and separation of Jupiter and Saturn as opposed to the gradual change in orbital locations from the traditional Nice model. The rapid speed and orbital changes induced from this “Jumping Jupiter” model could increase the chance for a smaller impactor to produce a larger disruption in the deep interiors of Jupiter and Saturn, as the collision velocity will have significantly increased. The second modified scenario is the “Grand Tack” model (Walsh et al. 2011), in which Jupiter forms at 3.5 AU, migrates to the interior of the Solar System at 1.5 AU, and subsequently migrates back to its current orbit of 5.2 AU. In this scenario, Jupiter traverses the asteroid belt not once, but *twice*, scattering the small planetesimals and certainly inducing many collisions with the Gas Giant. Here, as you have increased your amount of collisions, the likelihood for a single collision with the particular requirements to create a “fuzzy core” will increase, as well—now, only Saturn must beat the difficult odds of reaching this perfect collision, as well. Although both these scenarios may increase the likelihood of our Gas Giants evolving to host diffuse cores, the chances are still low, and we are as of yet unsure if either of these migrations truly occurred.

Conversely, it is exceptionally difficult to produce a “fuzzy” core from standard formation mechanisms, as convection is simply too efficient and consistently disrupts compositional gradients formed *via* a result of core accretion formation models with runaway accretion (whether this process is slow or rapid) in nearly all current formation modeling approaches

(Müller et al. 2020). The question, then, remains—what formation or development process creates planets with these diffuse, fuzzy cores? And should these planets be formed *via* core accretion, what can the “fuzziness” of the core teach us about the slow changes in rates of accretion of heavy elements to hydrogen and helium over time? (Helled & Stevenson 2017)

Despite all these questions and some of these answers, many questions still remain. The current gravitational measurements we have for Jupiter remain degenerate and non-invertible—we can produce models with vastly different interior structures, all of which accurately reproduce the observed gravity measurements. Because these gravitational measurements lose nearly all of their sensitivity (Helled et al. 2010) deep within the outer planets, the interior structure of Jupiter is rendered fully model-dependent. It is clear that new techniques must be applied to directly measure the inner half of the Jovian radius to complement *Juno*’s measurements and allow for further constraints on the deepest interior regions of the planet.

1.4 Dioseismology: Probing the Deep Jovian Interior with Seismology

It has been universally agreed upon by experts in the field that seismology of Jupiter—hereafter referred to as Dioseismology (Provost et al. 1994)—is the future for directly and unambiguously probing the deep internal structure of Jupiter (Private communication, *Juno* Prime Mission Results Workshop, 2022). Measurement of global modes propagating through the Jovian interior would provide valuable constraints for current Jovian interior models, helping break the degeneracies that are currently present from current non-invertible gravity

field measurements. This leap forward in knowledge, in turn, could reveal answers to the originally posed, burning question—by what process do Gas Giant planets form?

In addition to constraining the internal structure, detection of the Jovian global oscillations will also provide valuable information on properties of the oscillations, themselves. In particular, if these oscillations exist within Jupiter, what excitation mechanism is driving them? For the Sun, this excitation mechanism is the strong convection, but no such determination has been made regarding Jupiter—the excitation mechanism remains a mystery. Different excitation mechanisms will produce global modes with amplitudes and frequency regimes unique to the original excitation mechanism.

There are four main excitation mechanisms for the global modes of Jupiter currently considered by the community: convection (as in the Solar case), water storms (such as typical storms that we think of here on Earth—hurricanes—which we know to exist on Jupiter in the form of the Great Red Spot), impacts (for example, the Shoemaker-Levy 9 comet impactor in 1994), and deep-atmosphere silicate convection, manifesting as “rock storms”, the existence of which is currently only theoretical for the case of Jupiter. Current simulations, all as reported in Markham & Stevenson (2018), imply that oscillations excited by turbulent convection should have amplitudes of $\sim 0.1 \text{ cm s}^{-1}$, approximately three orders of magnitude smaller than those produced by a similar excitation mechanism in the Sun, and peaking at lower frequencies (lower than $1000 \mu\text{Hz}$). The expected amplitudes of oscillations driven by water storms are similar to that of turbulent convection—again, on the order of $\sim 0.1 \text{ cm s}^{-1}$, yet this time peaking at higher frequencies (beyond $2000 \mu\text{Hz}$).

For an impactor the size of the 1994 Shoemaker-Levy 9 comet, the largest body yet observed to collide with Jupiter, the resultant modes which could be excited and ring out throughout the interior of the planet are expected to have amplitudes on the order of \sim a few microns per second. (For this prediction, the assumption is made that an impactor of this size occurs roughly every 50 years.) Smaller impactors will produce even smaller amplitude oscillations; however, even the largest amplitudes of a few microns per second are so small as to effectively be negligible in the case of our current observational capabilities (Markham & Stevenson 2018).

Finally, we come to the case of deep silicate convection, or “rock storms”. These have are likely present within brown dwarfs, as models of deep interior storms accurately explain color variation of L dwarfs (Marley & Ackerman 2001), and predicted on hot exoplanets, and perhaps our own Jupiter, as posited by Fegley & Lodders (1994). However, these are purely theoretical for the case of Jupiter, and as of yet, have never been observed. Should these events exist, they will have much higher latent heat and energy to deposit into acoustic waves excited by their silicate convection, potentially large enough to produce amplitudes on the order of 10 cm s^{-1} , and peaking in the 1000–3000 μHz regime (Markham & Stevenson 2018).

It is clear that simply by producing a clear detection of Jovian global modes, and their corresponding frequencies and amplitudes, we can infer a great deal regarding the internal processes occurring within the planet. Should a detection occur with amplitudes higher than a few centimeters per second, the only feasible excitation mechanism to deposit that amount

of power is that of rock storms. Conversely, should no oscillatory signal be detected with amplitudes higher than this threshold, the existence of rock storms occurring within Jupiter and acting as an excitation mechanism may be entirely excluded. The pursuit of continuing the search for Jovian global oscillations is a difficult but noble one, as it holds the potential to shed light on many currently unanswered questions regarding the largest planet within our home Solar System.

1.4.1 Inheritance from Helioseismology

The field of Helioseismology inherits greatly from the widely successful parent field of helioseismology (Hart (1954), Leighton et al. (1962)), which measures acoustic standing waves, or pressure-modes that are stochastically generated by the energetic convection within the Sun, and subsequently propagate throughout its internal structure (Ulrich (1970), Leibacher & Stein (1971)). The speed of these waves will change as they propagate through varying densities, and the measurement of the waves—and subsequently, the sound-speed—at the photosphere of the Sun will contain information on the densities traveled through (Duvall et al. (1984), Christensen-Dalsgaard et al. (1985)). These waves are measured as a Doppler shift of Solar light, and the frequencies and amplitudes of these measured shifts can then be inverted to provide a sensitive and direct measurement of the internal Solar densities (Christensen-Dalsgaard 2002).

These historical efforts to make the sensitive measurements necessary for helioseismology utilized a Magneto-Optical Filter (MOF), or a sealed cell which contains the desired vapor of an absorption line that you wish to observe. This cell is subjected to a strong magnetic field

and a high applied temperature, inducing Zeeman splitting and Faraday rotation within the absorption line of interest. This results in two stable, sensitive wavelength probes (Cacciani & Fofi 1978), (Tomczyk et al. 1995), where a Doppler shift of the observed line will manifest as a change in intensity after passing through this MOF (a technique detailed further in Section 2.2.1).

Because these Solar oscillations are so small—with amplitudes of around 10 cm s^{-1} , proportional to an intensity change of just a few parts per million (Christensen-Dalsgaard 2002)—even on a bright, hot, and energetic body such as the Sun, this high instrumental sensitivity is imperative. It is also important to minimize external noise sources so that they do not overwhelm the signal; to include other velocity sources present on the Sun such as rotation and Solar activity. Most importantly, any periodic, temporal gaps in collected observations of these data will induce devastating side-lobes in the frequency analysis of these data—these side-lobes not only add a significant amount of excess noise, but decrease the amount of power in the central peak, as well.

To avoid these temporal gaps, while also collecting long-duration measurements (on the order of several months), many important discoveries in helioseismology have been made from multi-latitude telescope networks, such as the BiSON (Chaplin et al. 1996) and GONG arrays (Harvey et al. 1996). Space-based observations are excellent for helioseismic observations (such as the SOHO (Domingo et al. 1995), (Schou et al. 1998), and the SDO spacecraft (Pesnell et al. 2012)), as they have the potential to entirely avoid periodic observing gaps induced by the day/night cycle.

Finally, an interesting solution to avoiding periodic day/night gaps without the need for spacecraft is to collect observations from the geographic poles of the Earth. As the North Pole is an unsuitable location for astronomy, this leaves the South Pole—here, many breakthroughs have been made for helioseismology as well (Grec et al. (1980), Rolfe (1988), Duvall et al. (1991)), solidifying this geographical location as exceptionally valuable for seismic observations.

1.4.2 Historical Helioseismology Efforts

In historical searches for Jovian oscillations, Doppler velocimetry has often been considered as the most promising technique (e.g., Vorontsov et al. (1976), Schmider et al. (1991), Schmider et al. (2007), Gaulme et al. (2011)) due to its significant success for helioseismology. This technique is widely regarded as the best approach to track the atmospheric motions in the visible domain—these motions will be vertical and radial for seismic observations, and horizontal for wind circulation—and fundamentally consists of measurement of the Doppler shift of spectral Solar Fraunhofer lines which are reflected by the planet’s upper cloud layers. This is particularly advantageous, as the Doppler signal is enhanced by reflection (as described in Gaulme et al. (2018)).

For the field of planetary seismology, most attempts have been dedicated to Jupiter as it is the biggest and brightest planetary target in the outer Solar System as seen from Earth. However, these measurements are particularly difficult to conduct, for a multitude of reasons. Primarily, we are searching for very small velocimetric signals—at best, on the order of a few tens of centimeters per second—which will be overwhelmed by all other velocimetric

signals on Jupiter (e.g., atmospheric dynamics, upwelling, and in particular, the extremely rapid Jovian rotational signal of $\pm 12.6 \text{ km s}^{-1}$ at the equator). Additionally, any seismic signals will also be overwhelmed by velocimetric signals external to Jupiter, but induced by observations (e.g., the orbital velocity difference between Earth and Jupiter, and Earth's own rotational signal). For a more approachable reference as to the scale of these oscillations and the difficulty of collecting such a measurement: if we assume any oscillations present have an amplitude of $\sim 50 \text{ cm s}^{-1}$, this corresponds to attempting to measure the Jovian cloud decks moving approximately the length of a football field from 375 million miles away.

Finally, Jupiter is particularly dynamic, and many events which could be exciting oscillations (impact, water storm, and rock storms) may be variable, forming and dissipating, and thus, causing oscillations to ring out often, but not always. At first glance, this appears to introduce significant complications in the search for oscillations, as they may not be consistently excited. However, the low damping within Jupiter suggests that once any modes are excited, they are expected to ring out for up to 300 years at frequencies around $1000 \text{ } \mu\text{Hz}$ (Mosser 1995). The events which excite these oscillations (with the exception of impacts) should be occurring on timescales shorter than this, and thus, once modes are excited, they should be ringing out effectively consistently, subsequently remaining detectable.

Despite the complexities surrounding a potential search for these oscillations, they have been a prime observational target for Jovian scientists for nearly half a century. First observational attempts to search for the global oscillations with a MOF (Cacciani & Fofi 1978) were led by Schmider et al. (1991), then followed by observations with a Fourier-

transform spectrometer (Mosser et al. (1993), Mosser et al. (2000)), then, a double MOF (Cacciani et al. 2001). Most recently, observations have been collected with the first dedicated instrument, the Seismographic Imaging Interferometer for Monitoring of Planetary Atmospheres (SYMPA) (Schmider et al. (2007), Gaulme et al. (2008), Gaulme et al. (2011)), which is also a Fourier transform spectrometer.

Observations by Schmider et al. (1991), Mosser et al. (1993), Mosser et al. (2000), and Gaulme et al. (2011) all reported the presence of oscillations at a low Signal-to-Noise (S/N) level, with amplitudes between 10–100 cm s^{−1}. These previously reported values are in agreement with the maximum value allowed from theoretical works, which predict that the global oscillations of Jupiter should have amplitudes in the aforementioned regime of 10 to 100 cm s^{−1}, within a frequency range of 800–3500 μHz (Vorontsov et al. (1976), Bercovici & Schubert (1987)).

The most recent result—Gaulme et al. (2011)—reports that their detection, at an amplitude of 50 cm s^{−1} and at a frequency of 1200 μHz, likely represents the superposition of many modes, and the true amplitude of any individual mode may be a factor of two to three times lower. However, these reported amplitudes are perplexing, because an excitation mechanism that could create global modes with amplitudes of tens of cm s^{−1} would require an energy source to create and retain this energy that is *millions* of times more efficient for Jupiter than for the Sun (Markham & Stevenson 2018). Naturally, this is a bit concerning, as the Sun has a much higher mass and energy than is available on Jupiter. The only excitation mechanism capable of potentially producing modes with amplitudes of this level is deep,

energetic rock storms, which have not yet been confirmed to exist within Jupiter.

Additionally, all previous observations are collected from low-latitude sites, and subsequently, suffer from periodic gaps—and thus, spectral confusion from injected side-lobes. (Here, a pattern where the signal in any true peak within the Fourier transform is decreased, and this removed power is added in the form of separate, evenly spaced-noise peaks surrounding the true peak. This pattern, called a “window function”, is convolved with the entire signal, adding an overall comb-like structure and causing many different peaks to overlap). The effect of these side-lobes adds substantial complexity when searching for global modes, as this decreases the S/N even further. This low S/N in all previously reported detections has produced some uncertainty in the community about the existence of global modes at the reported amplitudes.

Regardless, the low S/N of previously reported detections certainly prohibits identification of any individual modes—this individual identification is necessary to perform any mode inversions to gain information about any boundary layers and internal structure within Jupiter. It is clear that further observations with a higher S/N are necessary to validate past reported detections.

Fortunately, work—aside from that contained within this dissertation—is also ongoing to help validate or counter past claims on detection of the global modes. In particular, the Jovian Interiors from Velocimetry Experiment (JIVE)/Jovian Oscillations through Velocity Images At several Longitudes (JOVIAL) experimental network, which uses a Mach-Zehnder interferometer to measure Jovian velocities, from atmospheric dynamics to global oscillations

(Schmider et al. (2019), Gonçalves et al. (2019)). This experiment is currently underway and implements a global approach to maximize S/N and reduce spectral confusion induced by the day/night cycle—by implementing a multi-latitude telescope network, with ground stations based in New Mexico, France, and Japan.

Naturally, this induces unique complications—in particular, synchronicity of observational time on three different telescopes (each with their own complexities and variations) in three different countries to ensure a minimum of observational gaps. As anyone who has tried to schedule observations during their desired time on even *one* telescope knows, the challenge of synchronizing *three* is a significant one, to be sure. However, successful observational results regarding atmospheric dynamics (Gonçalves et al. 2019), as expanded upon in Section 1.5.1 have already emerged, and this project holds promise for helping answer the standing questions regarding Jovian global oscillations and internal dynamics.

1.5 Jovian Atmospheric Science

Importantly, any measurement technique—including Doppler velocimetry—which can sense down to the deepest interior of the Gas Giants will also provide a wealth of further information on the atmospheric dynamics at play spanning from the surface-level, observable clouds to the convection, winds, and storms deeper in the interior of the planet.

1.5.1 The Jovian Zonal Wind Profile

In particular, these Doppler measurements can contain information regarding Jupiter’s Zonal Wind Profile (ZWP)—which is obtained by averaging the strong and stable winds circulating

Jupiter in longitude. This profile is considered to be a fundamental constraint of the Jovian atmosphere (Ingersoll et al. 2004)—and many historical observations have been dedicated to the study of this profile *via* cloud-tracking measurements (*Hubble* (Simon et al. 2015), *Voyager* (Gierasch et al. 1986), *Galileo* (Atkinson et al. 1998), ground-based (Hueso et al. 2017), etc.), all of which map the motions of features within the visible cloud decks as they progress throughout time.

However, an important consideration is that these cloud-tracking techniques fundamentally map the overall motion of the large-scale structures, and subsequently, they provide information on the velocity of iso-pressure regions instead of information regarding the true velocities of the particles which comprise the clouds. Measurements of the iso-pressure regions versus the true particle velocities could differ significantly if the cloud structures were subjected to thermal or gravity waves propagating through the Jovian atmosphere. Therefore, it would be beneficial to determine a separate measurement technique which holds the capability to confirm and validate that the cloud-tracking measurements are an accurate representation of the true particle velocities.

Thus far, the sole velocimetric measurements of Jupiter’s ZWP were performed with the dedicated instrument JOVIAL, which was inherited from SYMPA (Schmider et al. 2007) architecture by Gonçalves et al. (2019). It is important to note that, like the search for the global oscillations, this velocimetric measurement of the zonal flows surrounding Jupiter is exceedingly difficult to conduct—now, we are searching for clear velocities on the order of ~ 100 to 150 m s^{-1} (Johnson et al. 2018), which will again be overwhelmed and intermixed

with the rapid Jovian rotation of $\pm 12.6 \text{ km s}^{-1}$, so even the existence of this first Doppler measurement is a tremendous feat.

However, this single velocimetric measurement displays some significant differences from the well-studied cloud tracking ZWP. Aside from a lower resolution, the JOVIAL measurement of the ZWP is in direct opposition to the cloud-tracking measurement in the Northern Equatorial Belt (NEB), spanning the latitude region between 0–15 degrees. In particular, the JOVIAL measurement displays a minimum in velocity, while the cloud tracking profile displays a maximum here (Gonçalves et al. 2019). This difference holds a fundamental and substantial implication that, in the NEB regime, the true particle velocities may differ from the large-scale feature tracking velocities. A second velocimetric measurement is desired to either confirm or refute this result.

In addition to surface level measurements, the ZWP *beneath* the upper cloud decks has been studied through infrared temperature measurements. However, these measurements also rely on the structure characterized by the cloud-tracking measurements, and as such, are not an independent validation of the profile itself (Fletcher et al. 2016). Despite these extensive studies, many questions regarding the zonal winds still remain. What depth within the Jovian atmosphere do these winds extend to? Does the energy maintaining the winds arise from thermal convection from deep within the planet’s interior? Or are the winds caused by surface-level effects instead, such as variations in temperature between the alternating belts and zones, resulting in the winds being constrained to only a thin layer near the surface of the atmosphere? These questions have provoked curiosity for the past three decades

(Dowling 1995), and the search for the corresponding answers was one of the prime sources of interest driving the launch of the *Juno* spacecraft (Hubbard 1999).

1.5.2 The Odd Gravitational Moments as a Probe of Atmospheric Dynamics

Though the source and analysis of the even J-components has been discussed in detail earlier, their counterparts—the odd J-components—have not yet been addressed. These odd components would not exist if Jupiter had no dynamic effects aside from rotation—however, we know that this is not the case. Jupiter has strong winds encircling the planet, turbulent storms roiling within the clouds, and theoretically, global oscillations propagating throughout the interior of the planet.

Recent results from *Juno*’s gravimetric experiment revealed that the gravitational field includes a north-south asymmetry, the presence of which can only be explained as a result of atmospheric and internal dynamics (Kaspi et al. 2018). The odd harmonics J_3 to J_9 , from measurements collected from the *Juno* spacecraft, have confirmed that the ZWP extends to a depth of $\sim 3,000$ km, and suggest that there is a strong correlation between the surface-level clouds and this deep wind flow (Kaspi et al. 2018).

Although these gravimetry measurements are exquisitely sensitive to the outer locales of the planet, there exists a significant limit to their capabilities. In particular, inversion of these J-components is intrinsically non-unique—that is to say, many varying zonal wind profiles can produce the same observable gravimetric signals—and thus, any claims regarding the wind profiles and deep interior must be model dependent. In fact, *both* the models which suggest that internal flows and zonal winds are constrained to a thin, shallow level, and

models which report that the flows extend to the deep interior, are able to produce the odd J-component measurements collected by *Juno* (Kong et al. 2018).

Attaining a way to collect further direct measurements of these flows—as is possible with seismology—will intrinsically compliment the *Juno* gravimetric measurements of the winds, help progress understanding of how surface-level winds and deep winds are connected, and will help determine the origin and energy source of these strong, consistent winds.

Current work attempting to discover the source of the dynamic gravity reports that this gravity may be attributed to the global oscillations, implying that global modes with amplitudes ranging from 10–50 cm s^{−1} could produce the observed asymmetry, which previously could not be attributed to any other direct source (Durante et al. 2022). However, this field is particularly dynamic and rapidly evolving, and even more recent work (Militzer et al. 2022) implies that the suggest global modes are not, in fact, necessary to fit the observed asymmetric gravity field. Instead, they propose that a model containing a single set of reasonable assumptions, including a dilute core, a flexible model for the zonal wind speeds and their respective depths, a helium rain layer within the planet, and either no compact core, or a very small compact core at the center, can accurately reproduce all observed gravity harmonics, contradicting the report that global modes with an amplitude higher than 10 cm s^{−1} is necessary to produce the observed signal.

Further observations and constraints on the amplitudes of any global modes present within Jupiter, in addition to further measurements on the depths of the zonal winds, will help shed light on the processes which are truly occurring within the planet to produce the

observed gravity seen by the *Juno* spacecraft, helping to progress this dynamic and rapidly evolving field.

1.5.3 Applications of Doppler Velocimetry to Non-Jovian Atmospheric Science

Doppler velocimetry also holds potential for measuring atmospheric dynamics, not only of the Gas Giant planets but also of their orbiting moons—in particular, Io and Titan (Civeit et al. (2005), Luz et al. (2006)). It is also a valuable tool for studying the atmospheric dynamics of inner planets, such as Venus (e.g., Lellouch et al. (1994), Machado et al. (2017), Gaulme et al. (2019), and references contained therein).

Many historical Doppler spectroscopic studies of planetary bodies have been dedicated to Venus, in particular, to support the European Space Agency (ESA) Venus Express mission (Lellouch & Witasse 2008). These Venus observations were mostly performed by scanning across the planet with a single-fiber-fed, high-resolution spectrograph (e.g., Widemann et al. (2008), Machado et al. (2017), or through the use of long-slit spectrographs (Machado et al. (2012), Gaulme et al. (2019)). Finally, it is interesting to note that Doppler spectrometry has even been implemented to conduct wind velocity measurements of exoplanets (Louden & Wheatley (2015), Brogi et al. (2016)), suggesting that a wide array of future opportunities remain for scientific breakthroughs with this measurement technique.

1.6 Remaining Questions

It is clear that much remains to uncover in this turbulent and ever-changing field of Jovian atmospheric and interior science, and that further seismic observations are necessary to progress the field. These following questions are ones we hope to shed some light on throughout the course of this dissertation:

1.6.1 Remaining Questions Regarding Jovian Atmospheric Dynamics

Although we know much about the Jovian winds, there are other questions which remain aside from understanding the true particle velocities. In particular, how stable are these winds? On a long-term and global scale, this profile appears to be particularly stable, hosting velocities around $\sim 150 \pm 10 \text{ m s}^{-1}$). However, some local regions display variation on year-to-year timescales, which could benefit from further observations to strengthen current theories.

The NEB, defined as the band located in the region around $\sim 7^\circ \text{ N}$, displays a 4-year cycle cyclical variation. This region is known to be home to vertical upwelling, resulting in plumes and “hot spots” that are coupled alongside dark projections. These show the characteristic signatures of a trapped Equatorial Rossby wave (Barrado-Izagirre et al. 2013), (Arregi et al. 2006). Although surface-level, horizontal components of these features can be measured with feature-tracking techniques, determining the corresponding vertical components is more complex. Models report that hot spots such as these will likely develop vertical shear with a magnitude of $\sim 70 \text{ m s}^{-1}$ (Showman & Dowling 2000), which can be separated from the

larger stable flows of up to $\sim 150 \text{ m s}^{-1}$ as the larger flows are zonal and will thus have a maximum radial component near the limb of the disk, while the vertical shear is radial, and will thus have a maximum radial component near the center of the disk—as such, we can isolate the components of interest through observations on different regions of the Jovian disk.

In addition to the NEB variations, the Northern Temperate Belt (NTB), defined as the band located in the region around $\sim 21^\circ \text{ N}$, displays variation on a year-to-year scale, attributed to high-albedo plume outbreaks occurring every ~ 5 years (notably, the most recently observed outbreak occurred in 2016, implying that another outbreak should have occurred sometime in 2020). These plume outbreaks form strong features which encircle Jupiter (Sánchez-Lavega et al. 2017). The community has agreed that continued monitoring of these temporal changes and outbreaks, in addition to further advancement on modeling of these effects, holds potential to provide insight on the interplay between circulation on the outer regimes of the planet, and the periodic events which trigger them (Fletcher et al. 2020b).

Continued progress on constraining the vertical structure and temporal variability of Jupiter’s zonal winds is necessary to progress our understanding of the causes for and results of varying effects in the winds, and of the fundamental origin for the circulation persisting within Jupiter (Ingersoll & Pollard 1982). Through this work, we hope to answer specifically: *does the particle velocity of the zonal wind profile substantially differ from the overall cloud tracking profile? What are the true velocities present within these zonal flows?*

1.6.2 Remaining Questions Regarding Jovian Interior Structure

The first primary question regarding the Jovian interior structure pertains to the global modes themselves: *do global modes exist in the 1000–2000 μ Hz regime with amplitudes of 10–50 cm s⁻¹*, as reported by the SYMPA group (Gaulme et al. 2011) and explained as a solution to the dynamic gravity measured by *Juno* (Durante et al. 2022)? Or rather, does the lack of a repeated detection by the SYMPA group, and the capability to fully explain the dynamic gravity by allowing the internal structure and zonal flows to vary (Militzer et al. 2022), instead suggest that the global modes have not, as of yet, been detected?

If these modes do exist at the reported amplitudes, can we verify these detections, and if so, what can we learn about the Jovian interior structure from them? On the contrary, if the modes do not exist with amplitudes ranging from 10–50 cm s⁻¹, *what amplitudes do we expect these modes have?* If they are small enough as to be undetectable with current instrumentation and techniques, *what is the maximum possible amplitude they **could** have based on the minimum noise level of current techniques?* *What can these amplitudes, or constraints on amplitudes, tell us about the excitation mechanism(s) present deep within Jupiter?* Finally, if we are unable to detect the global modes with current technology, *what future progress must be made to measure the oscillations and finally place direct constraints on the Jovian interior?*

1.7 Dissertation Overview

We explore an attempt to progress the field of Jovian science through the design and development of a new set of sensitive, ground-based astronomical instrumentation, and a rich observational campaign collected through these instruments.

- Chapter Two details the concept, build, testing, and sensitivity of the instrumentation itself, along with its installation on the AEOS 3.6m telescope.
- Chapter Three discusses a twenty-four -night observational campaign implementing the newly developed instrumentation, along with detailing a novel data reduction and analysis pipeline, developed specifically for this instrumentation.
- Chapter Four discusses significant results from the new instrumentation and observational campaign, including a new velocimetric measurement of the Jovian zonal wind profile and comparison with historical profiles, the strictest constraints to date on the maximum possible amplitudes of the Jovian modes of oscillation, and inferences on new implications for allowed excitation mechanisms present within Jupiter.
- Chapter Five describes complications over the course of this experiment, and development of instrumentation to irrefutably validate our velocimetric sensitivity. Finally, we propose two separate solutions—one ground-based, and one space-based—to reach a higher sensitivity to improve future attempts to detect the Jovian global modes, in addition to detailing current and tangible progress towards making the ground-based path a reality in the near future.

- We conclude our findings in Chapter Six, which summarizes how our search for the Jovian oscillations has progressed the field and details the vast array of science which remains to be uncovered in the dataset collected through our novel instrumentation.

These sections precede an index of supplementary material, which includes the software developed in **MATLAB** over the course of this dissertation for reduction and analysis of our newly collected dataset.

CHAPTER 2

ON DEVELOPING INSTRUMENTATION FOR MEASUREMENT OF JOVIAN OSCILLATIONS

This chapter contains material published in Gulledge & Shaw (2022), Frontiers in Astronomy and Space Sciences, Astronomical Instrumentation Section.

From the questions posed in Chapter 1, it is clear that further observations of Jupiter with a more sensitive instrument are necessary to validate past detections on the global modes of Jupiter and to progress our understanding of not only structure which lies deep within the interior of the planet, but also the dynamic internal and atmospheric dynamics. To that end, we have developed a new and unique set of instruments: the Planetary Multilevel Oscillations and Dynamics Experiment (PMODE), which is built upon the experience and history of using MOFs for seismic observations of the Sun. In this chapter, we present the design and functionality of the PMODE instrumentation, along with its sensitivity and capabilities.

2.1 Instrumental Concept

Any radial motion of Jupiter’s observable cloud decks will result in a Doppler shift of light reflecting off of the particles—these shifts can be measured through the use of a Doppler imager. This type of instrument has been used to great success in helioseismology due to the underlying performance of the narrow pass-bands ($\sim 40 \text{ m}\text{\AA}$) which are created by the MOF system. The pass-band configuration produces a scenario where the high sensitivity and a wavelength stability of $0.0015 \text{ m}\text{\AA}$, or $\sim 6 \text{ cm s}^{-1}$ (Tomczyk et al. 1995) allow for

exceptionally precise velocimetric measurements. MOF-based approaches have proven to be extraordinarily successful in mapping out the interior structure of the Sun and other stars (asteroseismology). In terms of planetary seismology—again, Diaseismology, for the particular case of Jupiter—two extremely narrow pass-bands measure the shift of a Solar absorption line of interest, which is reflected from the Jovian clouds (Agnelli et al. (1975); Dick & Shay (1991); Tomczyk et al. (1995)), providing a sensitive measure of any radial motions of the clouds themselves.

We have implemented this powerful tool of Doppler velocimetry in a multi-channel instrument to measure the atmospheric dynamics and the global oscillations of Jupiter, called PMODE. Our primary instrument, the Doppler imager, is designed with two separate channels to measure small radial velocity shifts of both the 589 nm sodium doublet (which measures to a depth of approximately 3 bar within the Jovian atmosphere (Cacciani et al. 2001)) and the 770 nm potassium Solar Fraunhofer lines (measuring to a depth of approximately 0.7 bar (West et al. 2004)).

Figure 2.1 displays where these Fraunhofer wavelengths fall on the intrinsic full-disk Jovian albedo spectra with a 0.4 nm spectral resolution, as observed from the European Southern Observatory in July 1995 (Karkoschka 1998). Both the 589 nm sodium and 770 nm potassium channels probe regions of the Jovian spectra with relatively high albedo, which helps to maximize the overall light gain when inevitably restricting our filter size to the narrow widths of our MOF channels.

In addition to the Doppler imager, PMODE also hosts a polarimetric channel that is

designed to measure the linear polarization of light reflected off the Jovian clouds. This channel increases total spatial coverage on the Jovian disk (as changes in scattering angle will be the strongest closest to the edges of the planet, while radial Doppler motions are strongest nearest the center of the planet) while utilizing as much of the collected light for scientific purposes as possible, as opposed to simply discarding it. Simultaneous information on the Jovian atmosphere at levels probed by the 889 nm methane band (which is sensitive much higher in the Jovian atmosphere, to a depth of approximately 0.2 bar (West et al. 2004); this methane absorption band is also highlighted in Figure 2.1) may also prove to be a valuable diagnostic in the future. However, this channel of the instrumentation is not considered within the remainder of this dissertation, as we primarily dedicated our time and effort to the MOF based channels.

Implementing two separate channels within the Doppler imager provides the capability to concurrently probe two separate atmospheric levels of Jupiter, allowing for collection of three-dimensional measurements. In this dissertation, we focus on the full design of the instrumentation, but scientific validation and results are discussed specifically for the potassium channel of the Doppler imager side of the instrument, and particularly, for the 770 nm potassium channel, due to undesired detector artefacts within the 589 nm sodium channel, as will be detailed later in Section 2.4.

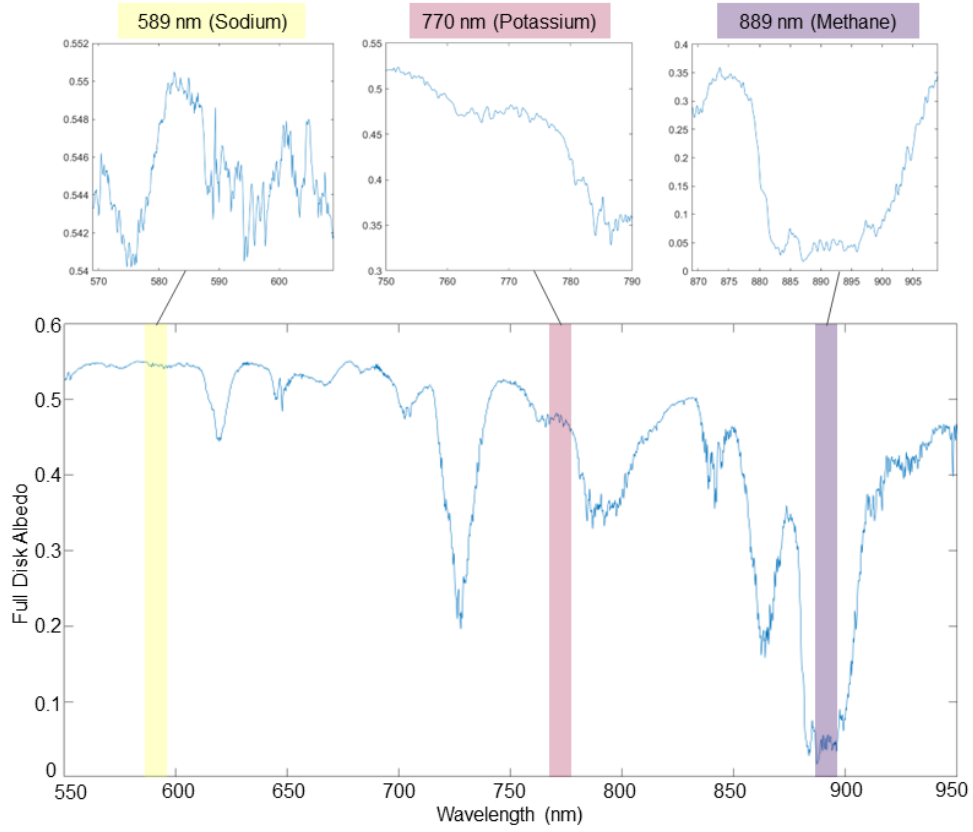


Figure 2.1 Intrinsic full-disk albedo spectra for Jupiter. Figure generated through data obtained through the NASA Planetary Data System, and originally collected from the European Southern Observatory in July 1995 (Karkoschka 1998). The bands in which we are observing with our MOF channels (yellow for Na, red for K) and polarimeter channel (purple) are highlighted, and a surrounding 40 nm zoom for each is inset surrounding the plot.

2.2 PMODE: the Doppler Imager Channels

Our Doppler velocimeter is designed to map the line-of-sight velocities in the troposphere of Jupiter. This is done *via* careful measurements of the Doppler shifts in reflected Solar Fraunhofer lines off of the Jovian cloud decks: in particular, the combined sodium D lines at 589.0 nm and 589.6 nm (unfortunately complicated by detector artifacts—however, we find it important to discuss the design of this channel regardless), and the potassium D₁ line at




Figure 2.2 Here, we show how light is affected as it travels through an atomic resonance vapor cell located between two crossed polarizers. Black arrows are representative of light from the wings of the absorption line of interest, while green arrows are representative of the remaining continuum light beyond the absorption line region. The light is first transmitted through a vertically oriented polarizer, and then into a heated glass cell (containing the vapor of the absorption line of interest), which is enclosed in a permanent rare-earth magnet assembly that applies a magnetic field parallel to the axis of light propagation. (1.) As a result of the temperature and magnetic field applied to the vapor cell, the light in the wings of the absorption line (black arrow) experiences Faraday rotation, which manifests as a rotation of the polarization state of the light. This rotation allows the light in the wings of the absorption line to pass through the Wollaston prism acting as a crossed polarizer following the vapor cell; this prism separates the light into two separate paths, one for each of the vertical and horizontal polarization states. These separated polarization states finally reach the detector, and the absorption line profile can now be seen as a dark band falling on the “MOF” image, which contains all our velocimetric sensitivity. (2.) The continuum light, outside the wings of the desired absorption line and represented by the green arrow, experiences no Faraday rotation within the MOF. Therefore, it travels through the system with no change to its original state, and transmits through the final Wollaston prism with the same vertical polarization state it began with. This produces a simultaneous “continuum” image on the detector.

770 nm. These two MOFs create narrow and sensitive pass-bands in the wings of the lines of interest; intensity changes in these pass-bands will directly correspond to Doppler velocities on the observed target.

2.2.1 MOF Optical Components and Transmission Path

Figure 2.2 shows a diagram which visualizes how a path of light is transmitted through a MOF, along with the physical effects induced on the light through this transmission. For our experiment, we implement MOFs which utilize a single-cell design (as detailed in Cacciani et al. (2001)) as opposed to the instruments used for helioseismology, where the Doppler measurement is obtained by differential measurement of the two wings in the line:

$$I = \frac{B - R}{B + R} \quad (2.1)$$

Here, I represents the measured intensity due to the absorption line shift, and B and R represent the intensities measured in the blue and red wings, respectively. Instead, PMODE uses only one line transmission profile, compared to the continuum, where our sensitivity to the Doppler shift is obtained in the wings of the spectral transmission profile:

$$I = \frac{B + R}{\text{continuum}} \quad (2.2)$$

In Figure 2.2, light passes through the assembly as follows:

- Light enters from the left, initially passing through a narrow filter (2 nm for the sodium channel, and 3 nm for the potassium channel) to eliminate undesired surrounding light.
- The incoming light then passes through a first polarizer, which allows only light matching its polarization state to be transmitted.
- This light then passes through a sealed glass cell, which contains either potassium or sodium vapor (depending on the channel of interest). A constant temperature is

applied to the vapor cell through an array of surrounding heaters, and a permanent rare-earth magnet assembly surrounding the cell applies a parallel magnetic field to the cell. This induces Faraday rotation of the light in the wings of the absorption line, resulting in an intrinsic change of the polarization state.

- After passing through the cell, the light is transmitted through another polarizing element (this time, it is a Wollaston prism), which is orthogonal to the first polarizing element prior to entry to the cell.
- At this Wollaston prism, the light is divided into two separate beams:
 - One beam passes through the assembly entirely unattenuated.
 - One beam is blocked by this final cross-polarizer. The only light which passes through is the narrow pass-band which has been rotated in polarization *via* the Faraday effect as it was transmitted through the MOF assembly.

Our output is two simultaneous beams, creating two separate images on the detector—the “MOF image” containing velocimetric information within the absorption line profile, and the 2–3 nm continuum image.

We divide this output MOF image by its corresponding continuum image to eliminate Jovian disk structure, produce a resultant measurement of the Doppler shift (as in Equation 2.2) in the form of a radial velocity map called a Dopplergram, which should be insensitive to albedo fluctuations (as any and all albedo fluctuations should be contained within the continuum

image as well, and thus, should cancel out upon division). Within this Dopplergram, any intensity changes should theoretically be due only to the Doppler shift of the light.

2.2.2 Doppler Velocity Spatial Signal on the Jovian Disk




Figure 2.3 Panel A: The “MOF” image, where the absorption line can be seen on the left portion of the disk. Panel B: The simultaneous “continuum” image. Panel C: The Dopplergram produced by division of the MOF image by the corresponding continuum image. Within this Dopplergram, pixel intensity values correspond to Doppler velocities towards and away from the observer within the dark absorption band. These frames were collected on the night of 12 August 2020 from the AEOS 3.6m telescope; here, a 13.76 km s^{-1} relative velocity between observer and Jovian disk center is present, and roughly 200 pixels span the entirety of the disk. Each of these images is produced from one single integration, with bias, dark, leak image, and flat-field calibration applied (these reduction steps are further detailed in Section 3.2). Intensities have been normalized within these frames for ease of viewing.

For our particular case, the instrument’s velocity signal (which can be seen in real data displayed in Panel A of Figure 2.3) appears as a projection of the absorption line of interest onto the Jovian disk. This differs from helioseismology, where the projection of the absorption line is not seen. This change occurs primarily as a result of different rotation rates—in particular, Jupiter has a very rapid rotation, with a period of 9h 55m, corresponding to $\pm 12.6 \text{ km s}^{-1}$ at the equator, resulting in a delta of 25.4 km s^{-1} across the disk. However,

the absorption lines we are using have widths significantly narrower than this. For the 770 nm potassium channel, our measured absorption line has a Full Width at Half Maximum (FWHM) of 0.164 Å, (as calculated from the National Institute of Standards and Technology (NIST) Atomic Spectra Lines Database (Kramida et al. 2021)). From the classic equation for Doppler shift,

$$\frac{\Delta\lambda}{\lambda_0} = \frac{v}{c} \quad (2.3)$$

where $\Delta\lambda$ is your change in wavelength from the rest wavelength λ_0 (here, 769.896 nm, or 7698.96 Å), v is the Doppler shift of interest, and c is the speed of light, in km s⁻¹, or $\sim 3 \times 10^5$ km s⁻¹, we can convert this FWHM to a Doppler shift as well, which results in a FWHM of 6.386 km s⁻¹. A visualization of the 770 nm potassium absorption line, as a function of both wavelength and Doppler shift, is displayed in Figure 2.4. Because this range over Doppler velocities that the absorption line is sensitive to is significantly smaller than the 25.4 km s⁻¹ velocity delta across the Jovian disk, a substantial portion of the disk is Doppler shifted out of the range of spectral sensitivity that our MOF can measure.

Figure 2.6A displays how the rotation manifests in a 770 nm potassium MOF image. Here, the Solar line (following convolution with the filter passband) is much narrower than the overall velocimetric change induced by the Jovian rotation. Because of this effect, the spectral line is resolved within the final image present on the detector, and manifests as a dark, vertical band on the Jovian disk. The steepness of this line corresponds to the sensitivity of the MOF, so this particularly narrow, steep absorption line provides high local sensitivity. However, this results in little-to-no velocity sensitivity on the majority of the




Figure 2.4 Visualization of the 770 nm potassium absorption line with normalized intensity, as a function of both wavelength (bottom axis) and Doppler shift (top axis). Data obtained from the NIST Atomic Spectra Lines Database.

Jovian disk, as it is Doppler shifted beyond the range that the absorption line is capable of sensing.

Importantly, this disk limitation exacerbates spatial aliasing in our measurements, as we are searching for modes with a global signal (which can be represented by spherical harmonics) across the entire planet, and we have a severely limited regime of the planet on which we can directly measure any signal. To explain this another way—if we were looking to measure a mode represented by the $Y_\ell^m = (1, 0)$ spherical harmonic—and we had the capability to view and measure the entire surface of Jupiter at once (an impossibility, as at any given time, we can only see one-half of the planet)—we would measure *only* the

$Y_\ell^m = (1, 0)$ mode and no others, as spherical harmonics are orthogonal over an entire sphere.

However, if we restrict our viewing to a smaller region of the disk—now, for example, the dark band displayed in Figure 2.6A—and attempt to measure the same $Y_\ell^m = (1, 0)$ mode, we will measure not only the desired mode, but also some signals from the surrounding modes (e.g., perhaps the $Y_\ell^m = (1, 1)$ or $Y^{-\ell^m} = (1, -1)$ modes). This adds a good degree of difficulty when attempting to identify specific and individual modes, a complication which is analyzed more thoroughly in Section 3.6.2.

In the PMODE Doppler imager, we attempt to mitigate this problem by adding a second MOF channel which is sensitive to the Na D lines at 589 nm. Importantly, these lines are much broader (which can be seen in Figure 2.5), maintaining a FWHM of ~ 0.436 nm or 22.192 Å.

The increased broadness of the Na lines extends our region of velocimetric sensitivity to nearly the limb of the planet (as shown in Figure 2.6B), and following combination with the K 770 nm line, our two Doppler channels should provide high sensitivity over nearly all of the visible portion of Jupiter, as shown in Figure 2.6C.

An example of output images collected with this instrument package on the Advanced Electro Optical System (AEOS) 3.6m telescope over the course of a 24-night observing run (detailed in Chapter 3) and processed through a novel data reduction pipeline written in MATLAB (which applies bias-, dark- and flat-field corrections, further detailed in Section 3.2, as well) are shown in Figure 2.3.




Figure 2.5 Visualization of the 889 nm Sodium DI absorption line with normalized intensity, as a function of both wavelength (bottom axis) and Doppler shift (top axis). Data obtained from the NIST Atomic Spectra Lines Database.

2.3 Full Experiment Optical Design

PMODE was designed to have the capability for future integration on the Cassegrain port of a 0.5m Ritchey–Chrétien telescope; the reasoning behind this choice will be expanded upon in Section 5.3.1.

However, for use on a breadboard in a Coudé room of a larger telescope (in particular, the AEOS 3.6m telescope), the simple optical layout with a small footprint (see Figure 2.7) can be expanded to cover more area, while maintaining the general overall structure and optical layout. The primary modification between these designs manifested in a larger initial de-magnification of the input light from the AEOS 3.6m telescope, the beam of which is




Figure 2.6 Simulation displaying sensitivity of the MOF channels of PMODE. Here, a lower value represents a higher sensitivity, as we are displaying the ability to measure smaller Doppler shifts—therefore, a sensitivity of 200 m s^{-1} is better than a sensitivity of 2000 m s^{-1} . The 770 nm potassium MOF produces a narrow and steep line on the Jovian disk (Panel A). Meanwhile, the 589 nm sodium MOF produces a passband which is broader, covering a larger area on the disk, but is also less sensitive (Panel B). The combination of these two cells (Panel C) displays the theoretical full sensitivity and spatial coverage, spanning the majority of the Jovian disk.

naturally larger than that of a 0.5 m telescope.

Figure 2.8 shows the overall optical path for PMODE. Prior to entry into our optical components, the native 60 arcsecond field-of-view enters into the facility Adaptive Optics (AO) system (Lewis C. Roberts & Neyman 2002), within which we have installed an uncoated window that produces a Fresnel reflection of our target, to then be used for tip-tilt correction on significantly extended targets such as Jupiter, and additional higher-order wavefront measurement and full correction when observing much smaller targets, such as Uranus and Venus.

Following transmission through the AO system, the collimated light enters the Coudé room, where an initial 4:1 SORL off-axis-parabola beam reducer compresses the beam size so that it will fit on subsequent 1-inch optics throughout the system. Following beam reduction




Figure 2.7 Left panel (A): A Zemax model of small-footprint original optical design for PMODE. Middle panel (B): a rendered Computer Aided Design (CAD) drawing of the originally designed instrument suite. The left side of the path contains the polarimeter, tracker, and wavefront sensor, none of which are considered within this dissertation. The right side of the path (blue and green lines) contains the sodium and potassium MOF channels. Right panel (C): The constructed system as of July 2019; the innermost channel is missing from the Doppler imager. No photographs were collected of the completed PMODE instrument, so we provide these of the smaller breadboard footprint, for reference.

via the SORL, the light then encounters OE (corresponding to the leftmost column in the table from Figure 2.8) #1, an F500 achromat lens (and the first part of a pupil relay system).

A quick refresher on the pupil plane vs. the image plane is shown in Figure 2.9—the pupil image is a re-image of the entrance aperture, or primary mirror, of your telescope. All fields overlap at this pupil image, and it contains all information on the phase and amplitude of aberrating wavefronts for your image. In this pupil image, your light is collimated, neither converging nor diverging.

OE1 causes the light to converge. This converging light then reflects off a fold mirror (OE2), coming to a focus at OE3, which is an iris with a small opening, acting as a field stop to let only the focused light through and blocking any undesired surrounding light. The

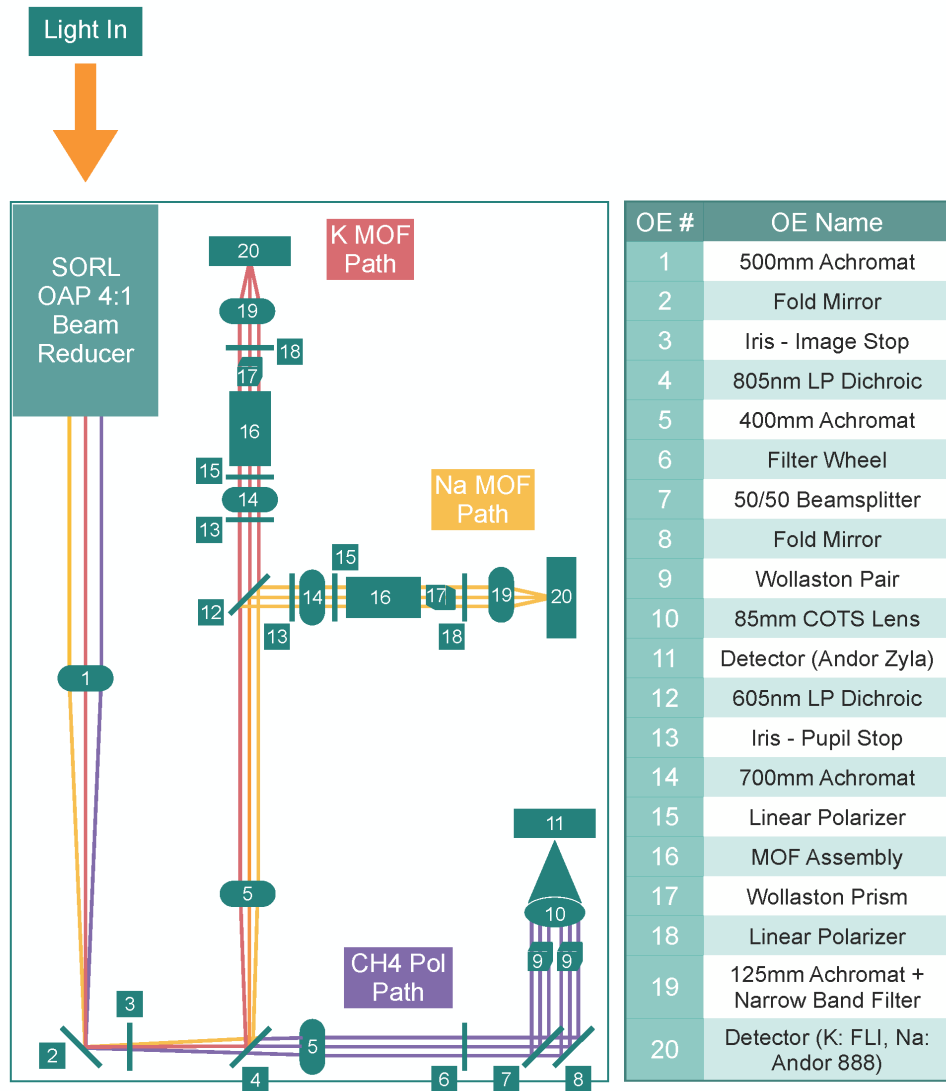


Figure 2.8 Diagram of the PMODE optical path following its installment in a Coudé room of the AEOS 3.6m telescope. On the left is a birds-eye view of the optical bench; the table on the right provides the corresponding information for each Optical Element (OE) that corresponds to the numbered component on the bench. Incoming light enters a beam reducer in the top-left of the bench, and is subsequently split by dichroic long-pass filters into three separate paths, each displayed as a representative color on the diagram. The first dichroic transmits all wavelengths beyond $\lambda = 805$ nm, creating a path for the 889 nm polarimeter channel. This path is shaded purple. The reflected light from the first dichroic is later split into the red path (which represents the second longest wavelengths and creates the 770 nm potassium MOF channel) and the yellow path represents the 589 nm sodium MOF channel.

light then passes through OE4, a $\lambda=805$ nm long-pass dichroic filter. This thin-film, coated filter transmits all light longer than $\lambda=805$ nm, and reflects all light shorter than $\lambda=805$ nm.

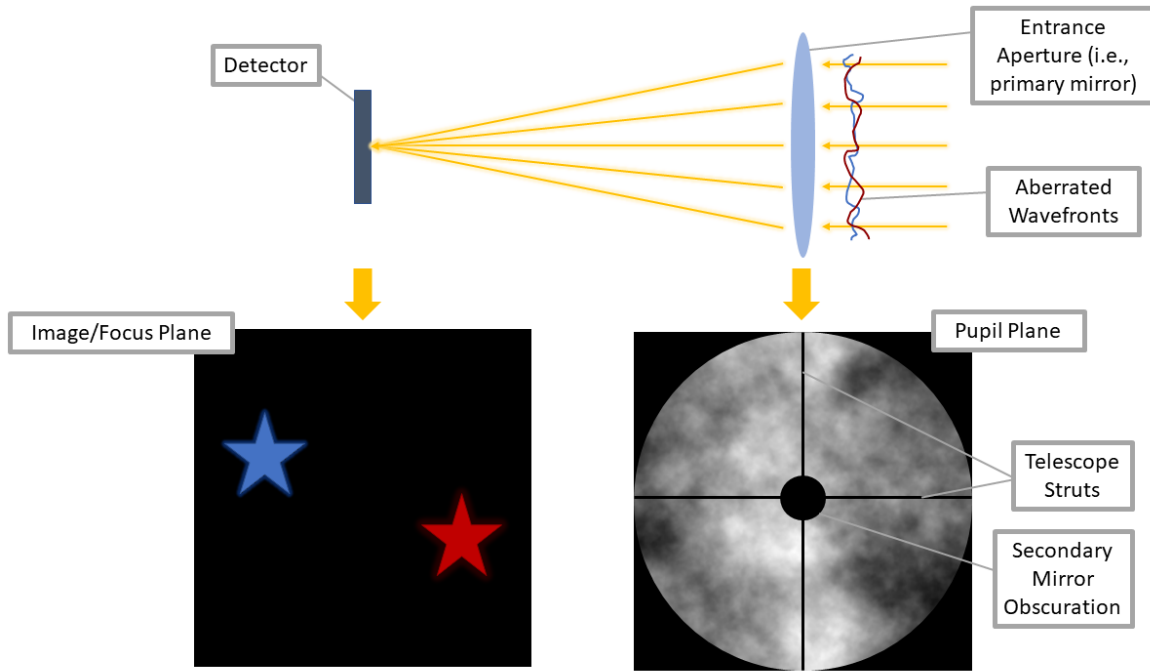


Figure 2.9 The image plane (left) vs. the pupil plane (right) for an optical system. The image plane displays your targets in focus, and the pupil plane—an image of your entrance aperture, likely the primary mirror of your telescope—contains information about the wavefront of your targets

2.3.1 PMODE Full Optical Design—889 nm Polarimeter Path

We will first follow the transmitted light, all wavelengths longer than $\lambda=805$ nm. This contains the polarimeter arm of our instrument, which is not considered in the analysis of this dissertation, but is included in this optical path description for completeness. This beam follows the bottom edge of the diagram shown in Figure 2.8. After transmission through the dichroic, the light then passes through OE5, an F400 achromat lens, completing the pupil relay system initialized at OE1. OE5 re-collimates the light, so that transmitted light beyond this optical element is in a pupil plane. Next, the light passes through OE6, a filter wheel

containing the Sloan Digital Sky Survey (SDSS) r and z filters, an open element that allows all incident light to pass, and an 889 nm filter, isolating specifically the light corresponding to the desired Jovian methane absorption band which the polarimeter is designed to sense.

Following the filter wheel, the light then encounters OE7, a 50/50 beamsplitter, which transmits 50% of the light, and reflects 50% of the light. The transmitted portion of light through OE7 then bounces off OE8, which is a fold mirror to send the light directly up, parallel to the reflected path from OE7. Each of these paths (reflected and transmitted) now passes through its respective Wollaston prism (an optical element which is made of two right-triangle calcite components, which separates incident incoming light into two separate and linearly polarized output beams, with polarization states orthogonal to each other), a pair of which is shown as OE9 in Figure 2.8.

After separation into these four separate polarized states by OE9, the light then passes into OE10, which is an 85-mm off-the-shelf photography lens to focus the incoming collimated light onto OE11, which is an Andor Zyla sCMOS detector to record the light for the polarimeter path, producing four final images with four separate polarization states onto this detector.

2.3.2 PMODE Full Optical Design—589 nm Sodium Path

Returning to OE4, the $\lambda=805$ nm dichroic, we can now begin following the reflected path, which follows the upwards direction on Figure 2.8. This light also first encounters another OE5 component, or an F400 achromat lens. As in the polarimeter path, this lens completes the pupil relay system initialized at OE1, so that transmitted light beyond this optical

element is collimated and in a pupil plane. Next, the collimated light encounters OE12, another long-pass dichroic, this time separating the light at a wavelength of $\lambda=605$. Light longer than this wavelength is transmitted for the potassium MOF channel, a path which we will follow in Section 2.3.3.

For now, we instead follow the reflected path for the sodium MOF channel, which sends the light towards the right direction in Figure 2.8. This light encounters OE13, which is a second iris to act as a field stop, excluding all light beyond the rough size of Jupiter's diameter. This light next encounters OE14, a clean-up 700 mm achromat lens, which ensures that the light is fully collimated prior to entry into the MOF assembly.

This collimated light next passes through OE15, a linear polarizer which is the first element in the full MOF assembly as described in Section 2.2.1 and represented in Figure 2.2, and as such, we will only give a brief description here. OE15 allows only light matching a defined input polarization to enter into the Na MOF vapor cell, represented as OE16 in Figure 2.8, which again, applies temperature and a magnetic field to the vapor, inducing Faraday rotation in the wings of the absorption line, and allowing only this rotated light to pass through the crossed-polarizer at the output of the MOF, or OE17.

Here, OE17 is another Wollaston prism, which splits the light into two orthogonal polarization states, one matching the input state (and producing the continuum image on the detector), and the other allowing only the rotated light to transmit, producing the tracing of the absorption line on the Jovian disk. (The inclusion of this Wollaston prism at the exit of the MOF as opposed to a second linear polarizer was a modification to the original design,

and further detailed in 2.3.4.)

Next, the light encounters OE18, which is an additional linear polarizer, designed to decrease the intensity in the continuum image so that is closer in intensity to our (significantly dimmer) MOF image, to allow for longer exposures without saturating our detectors so that we can maximize S/N in our velocity sensitive image.

Finally, the light then passes through OE19, which consists of both a 125 mm achromat lens to focus the light onto the detector, and also of a narrow-band, 3 nm filter, which restricts the continuum image to the narrow spectral range immediately surrounding the absorption line so that we are sensing the same region and structure of the Jovian atmosphere. This light concludes its journey as two separate images—the MOF image and the continuum image—on OE20, an Andor iXon Ultra 888 Electron Multiplying Charge-Coupled Device (EMCCD) detector for this Sodium channel (yet another modification from the original design, also detailed further in Section 2.3.4).

2.3.3 PMODE Full Optical Design—770 nm Potassium Path

We return now to OE12, the second long-pass dichroic in our optical system, which, again, separates the light at a wavelength of $\lambda=605$. Now, we follow the transmitted path (which sends the light towards the upward direction in Figure 2.8), for the potassium MOF channel. This light path is very similar in both purpose and components to the potassium path (Section 2.3.2), so we can journey through this path at a more rapid pace.

Following transmission through the dichroic, this light encounters its own OE13, which again, is a second iris to act as a field stop, excluding all light beyond Jupiter’s diameter.

Next, the light encounters OE14, OE15, OE16, and OE17, which are the clean-up collimating lens, linear polarizer, potassium vapor cell, and Wollaston prism respectively, to make up the components for the potassium MOF assembly (again, as in Figure 2.2).

After transmission through the MOF, we once more have two images with orthogonal polarization states, which then encounter OE18, the final linear polarizer to decrease continuum image intensity, and then OE19, which consists of both the 125 mm achromat to focus the pupil image into an image plane on the detector, and 2 nm narrow-band filter to define the continuum image so that the only transmitted light is that which surrounds the absorption line of interest.

Finally, the light from the potassium path concludes its journey on OE20, the detector to record the light from both the MOF and continuum image simultaneously. For the 770 nm potassium channel, this detector was a Finger Lakes Instrumentation (FLI) Microline.

2.3.4 PMODE Full Optical Design—Deviations from the Original Design

The original plan for our MOF channels was to implement dual Wollaston prisms surrounding the MOF assembly, because these Wollaston prisms provide a higher extinction ratio than typical linear polarizers. However, the separation of light into two beams prior to entrance into the vapor cell resulted in a divergence of light which was too large, causing the light to expand rapidly and to overfill the entrance aperture for the vapor cell, inducing scattered light off of the glass cell walls.

Exchanging the entrance Wollaston for single-beam linear polarizers (which have an extinction ratio of 10^{-5}) avoids this divergence. Implementation of these polarizers in place of

the Wollaston prisms results in an overall $\sqrt{2}$ loss in the total Signal-to-Noise Ratio (SNR), but was a necessary modification to remove this scattered light throughout the vapor cells (which, if left uncorrected, could manifest as a false apparent velocimetric signal).

Finally, we added approximately 4° of tilt to the Wollaston prisms to reduce any optical distortion that could induce a variation between the continuum image and the MOF image. These two must be as identical in structure as can be feasibly achieved in order to retrieve the true line-of-sight velocity signal. This tilting technique reduces the relative difference between the two optical paths emerging from the Wollaston prism, a technique which is described further in Simon (1986).

The swapping of this input Wollaston to a linear polarizer also necessitated the final polarizer placed after the Wollaston prism immediately following the vapor cells, as the output continuum image was now significantly brighter than the MOF image. This clean-up polarizer was empirically adjusted to maximize the MOF signal, and is nearly crossed with the continuum image to drastically decrease its flux, so that the two images are nearly equal in intensity. The optical path shown in Figure 2.8 includes all of these modifications.

The initial optical design for the system implemented FLI Microline cameras for both the 589 nm sodium and 770 nm potassium MOF channels, for consistency across both arms of the experiment. However, when setting up and testing the instrumentation immediately prior to awarded observational time, the FLI Microline camera for the sodium channel inexplicably failed, and would often refuse to turn on whatsoever, or occasionally flicker between its “on” and “off” state. Naturally, for a long-term experiment on the search for small-amplitude

periodic signals, a camera performing in this manner is unusable. Fortunately, a backup camera was located at the observatory—in particular, an Andor iXon 888 EMCCD detector.

Unfortunately, as this occurred immediately prior to the week of the observational campaign, we did not have time to thoroughly test and characterize this detector; although, at first glance, its performance appeared to be more than adequate. However, due to the lack of thorough initial testing and characterization, this replacement camera ultimately exhibited hidden detector artefacts significant enough to prevent time-domain analysis of any data collected through the sodium channel.

2.4 Sodium Channel Detector Artefacts

The Andor iXon 888 EMCCD detector which replaced the original FLI detector for the Sodium channel contains a thinned, back-illuminated Charge-Coupled Device (CCD). Although this type of CCD provides a higher quantum efficiency than front-illuminated CCDs, at longer wavelengths, these detectors occasionally produce etalon fringes, or constructive and destructive interference of the light reflecting off of the silicon surfaces and interacting with itself. Andor, the company which produces these cameras, states that this highly disruptive effect can modulate incident signal by as much as 40%. Although this complication is known to exist, it typically does not appear for wavelengths of light below 700 nm; thus, we assumed that use of this detector for our 589 nm sodium channel would provide no substantial complications.

Unfortunately, this assumption was incorrect, and the etaloning fringes are not only

present within the Sodium data, but they are immensely complicating for future analysis. Firstly, because these artefacts remained undetected until the end of the observational campaign—the fringes are invisible within the continuum and velocity images (Figures 2.10A and 2.10B), becoming apparent to the eye only once Dopplergram division to remove the Jovian disk structure is completed (Figure 2.10C).

Secondly, these artefacts are time-varying in intensity—manifesting similarly to the effect of any true oscillatory signal within these data. Thirdly, the artefacts are not only time-varying, but are spatially varying as well (as shown by the changing structure of the fringes from Figure 2.10C to Figure 2.10D). This prevents any simple averaging or subtraction to remove them from the dataset.

Combined, these complications are intrinsically tangled within the dataset. The spatial width of features in the added noise signal are on par with zonal wind features which we hope to measure within the Jovian atmosphere, preventing not only the search for global modes within these data, but also the search for atmospheric dynamic effects as well.

We have attempted to remove these artefacts from these data in a variety of ways: median-combining all frames in the hopes that the fringes average out for atmospheric dynamic analysis (they do not), fitting curves to the fringes through a series of vertical cuts for subsequent subtraction (however, this also fits and removes any zonal wind signatures, and also results in subtracting varying, individually-fitted structure, a certain complication for searching for small, time varying signals that will be changing these data), and Fourier filtering (however, the absorption line profile is also of comparable width, and this filtering

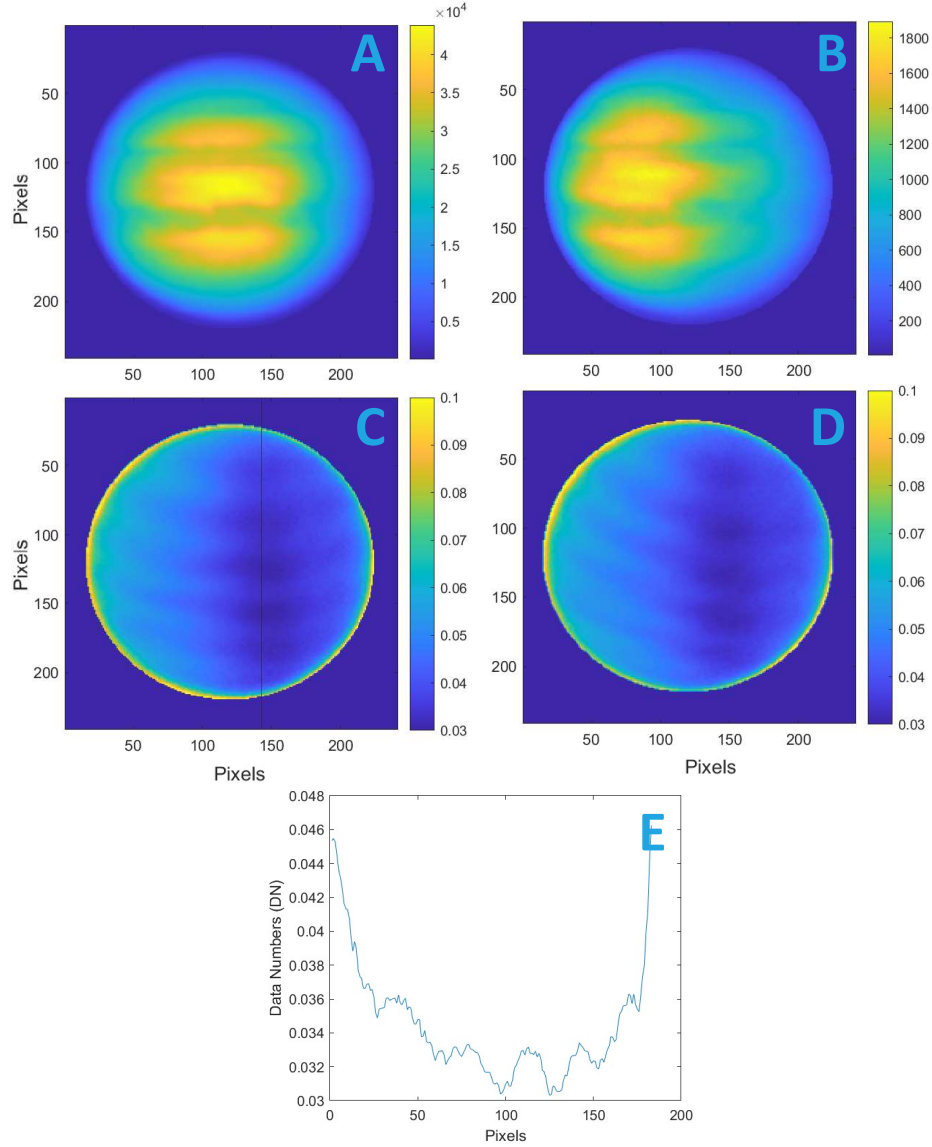


Figure 2.10 Panel A: Example continuum frame for the Na data. Panel B: Example MOF frame for the Na data. Panel C: Example Dopplergram from the Na data (created from Panels A and B), from early in the night of 15 August 2020. Panel D: A second Dopplergram from 15 August 2020, collected later in the night. Both Dopplergrams display the fringes, varying in structure through the night. Panel E: A projected cut along the vertical black line in Panel C, displaying how strongly these etaloning fringes affect our data.

removes the sensitive line and its' subsequent changes entirely). As all attempts thus far to disentangle the noise have been in vain, we choose instead to remove the sodium dataset from the scope of our analysis throughout the remainder of this work and focus solely on the potassium channel of our Doppler velocimeter.

2.5 Theoretical Performance

As with any instrumentation, it is valuable not only to understand how your experiment truly performs, but also to understand its theoretical capabilities. To evaluate this, we created a simple model to characterize the sensitivity for the Doppler velocimeter—specifically, for the potassium channel, as the sodium channel was plagued by the aforementioned detector artefacts. The goal of this calculation is to estimate the photon noise level of an observation of Jupiter collected on the 770 nm potassium FLI Microline detector. All procedures and equations are adapted from techniques detailed by Gaulme et al. (2011), Gaulme et al. (2008), and private communication between Dr. Gaulme and our collaborators.

Because our velocity sensitivity is directly represented as a change in intensity within our images, we first must determine the equation which will describe the change in the recorded flux on the detector. This can be represented as the derivative of the absorption line of interest (because the sensitivity is directly influenced by the slope of the absorption line—a narrower line with a steeper slope produces a higher sensitivity), convolved with the dual-probe pass-bands produced by the MOF, as follows:

$$S = \frac{\frac{dF}{dv} * I}{F * I} \quad (2.4)$$

Here, S represents your final sensitivity, F represents your photon flux measured on the detector, and I represents overall intensity as a function of your Doppler velocity, v).

As has been hinted at a few times thus far, the steepness of the absorption line (and subsequently, its derivative) is one of the largest factors contributing to the effectiveness of MOF-based techniques, as the sensitivity of your instrument to intensity directly changes with the steepness of the line.

For rapidly rotating targets of interest, such as Jupiter, choosing a line which has a narrow and steep profile creates higher local sensitivity, but restricts the area on the disk on which you have any sensitivity to conduct a measurement (an effect examined further in Section 3.6, and the reason why we expect a higher sensitivity for our potassium channel).

Wider Solar lines (such as the sodium channel) are shallower and less sensitive, but cover a larger area on the Jovian disk. The visual width of these lines on the disk is determined by the rotation rate of the target, and how much of the spatial area is Doppler shifted out of the width of the line, as detailed in Section 2.2.2; therefore, for targets which are rotating more slowly (such as the Sun), or targets that are unresolved, the behavior of the lines on the target will differ.

For the purposes of this analysis, we make the simplifying assumption that Jupiter is a perfectly spherical body with solid-body rotation, and ignore the effects of differential rotation, an assumption that we believe to be acceptable, as the effects from differential rotation are on the order of 1% of the average Jovian rotation period (private communication between P. Gaulme and our collaborators).

Next, to determine the flux from Jupiter that reaches the Earth, we begin with the amount of Solar flux that reaches the Earth at our wavelength of interest, 770 nm, which is:

$$F_{\odot To \oplus} = 4.65 \times 10^{18} h\nu/s/nm/m^2 \quad (2.5)$$

Because the Earth is at 1 Astronomical Unit (AU), Jupiter is at ~ 5.2 AU, and flux falls off in a predictable way ($\propto 1/r^2$), we can expand this value to the Solar flux reaching Jupiter as:

$$F_{\odot To \Jup} = F_{\odot To \oplus}/5.2^2 = 1.72 \times 10^{17} h\nu/s/nm/m^2 \quad (2.6)$$

Combining this value with the radius of Jupiter ($R_{\Jup} = 7.1492 \times 10^7$ m) gives us the total integrated Solar flux arriving at Jupiter (at a wavelength of 770 nm), as follows:

$$F_{\odot To \Jup Tot} = F_{\odot To \Jup} \pi R_{\Jup}^2 = (1.72e17 h\nu/s/nm/m^2) (7.145e7 m)^2 \pi = 2.76 \times 10^{33} h\nu/s/nm \quad (2.7)$$

From this total flux, we now calculate the reflected flux off the Jovian clouds. The Jovian albedo at 770 nm is $A_{\Jup} = 0.46$ (Karkoschka 1998), and multiplying this albedo value into the total received flux, $F_{\odot To \Jup Tot}$, gives us the reflected value, as:

$$F_{\Jup Reflect} = A_{\Jup} \times F_{\odot To \Jup Tot} = 0.46 \times 2.76 \times 10^{33} h\nu/s/nm = 1.2696 \times 10^{33} h\nu/s/nm \quad (2.8)$$

Finally, to obtain the flux value we will use for our calculation, we now determine the amount of flux $F(r)$ that we receive at a distance of r from Jupiter:

$$F(r) = \frac{1.2696 \times 10^{33}}{4\pi r^2} h\nu/s/nm \quad (2.9)$$

To determine our theoretical best performance, we will assume that Jupiter is at opposition,

Optical Element	Transmission
Prior to Coudé	0.5
OAP 1	0.99
OAP 2	0.99
Achromat	0.99
Fold Mirror	0.99
Dichroic	0.98
Achromat	0.99
Dichroic	0.98
Achromat	0.99
Polarizer	0.425
MOF Assembly	0.83
Wollaston	0.425
Polarizer	0.425
Achromat	0.99
NB Filter	0.5
Total: 0.042	

Table 2.1 Transmission of optical components for a MOF arm within PMODE. Approximately 4% of collected light from the telescope reaches each of our MOF detectors following transmission through the respective arm.

or at a distance of roughly $6.3 \times 10^{11} \text{m}$, for a received flux of:

$$F(r) = \frac{1.2696 \times 10^{33}}{4\pi (6.3 \times 10^{11} \text{m})^2} h\nu/s/nm = 2.55 \times 10^8 h\nu/s/nm/m^2 \quad (2.10)$$

The final transmission of light through all optical components in one arm of the PMODE instrument (detailed in Table 2.1) is approximately 4.2%. The quantum efficiency of the characterized 770 nm FLI Microline detector is approximately 70%, and the equivalent width of our potassium absorption line is 0.164 Å, or 0.0164 nm. Finally, the diameter of the primary mirror of the AEOS telescope is 3.67m. Combining these values with Equation 2.10 gives us the total number of photo-electrons measured by the detector each second:

$$N_{Phot_{\gamma}} = 2.55 \times 10^8 h\nu/s/nm/m^2 \times 0.042 \times 0.7 \times \pi (3.67/2\text{m}) \times (0.0164\text{nm}) h\nu/s \quad (2.11)$$

This produces a final value of 6.92×10^5 photo-electrons incident on the detector per second.

For a 28-second exposure, this becomes 1.94×10^7 photo-electrons incident on the detector per

exposure. Coupling this incident flux with resolution and seeing effect, and limb-darkening allows us to model Jupiter *via* MATLAB.

For the velocimetric sensitivity, we combine the Jovian rotation rate and the line of sight velocity to determine the overall radial velocity sensitivity; the same technique that is used to obtain our “MOF Unit” for calibration (which is detailed in Equation 3.6). With this, the expected instrumental velocity sensitivity (S_I) at each pixel can be calculated as:

$$S_I = \frac{1}{S_I \sqrt{N_{Phot} \eta_4}}. \quad (2.12)$$

This calculation produces the results shown in Figure 2.6. This simulation was created from similar techniques detailed in Gaulme et al. (2011), and from further private communication from the author with the PMODE team.

This simulation produces a theoretical value for the lowest attainable noise level for the PMODE instrument of $\sim 240 \text{ m s}^{-1}$ per pixel, per 28-second exposure. If we exclude regions nearer the limb of Jupiter that have higher noise levels (which will also be excluded in final analysis), and consider only pixels with respective noise levels below 1000 m s^{-1} , we obtain an average sensitivity of 616 m s^{-1} . This sensitivity increases still further, as we integrate the intensity in each image to produce our final data product (detailed further in Section 3.2). The summation of all the sensitive pixels produces a final theoretical value of a $\sim 7 \text{ m s}^{-1}$ attained sensitivity per 28-second exposure.

On-sky performance of the instrument is comparable to this theoretical estimate. Following data reduction and final on-sky calibration (all thoroughly detailed in Sections 3.2 and 3.0.2), our 770 nm potassium MOF channel images have a plate scale of 0.22 arcsec-

onds per pixel, oversampling Jupiter by a factor of approximately 3 when considering seeing effects. Within one 28-second exposure, the standard deviation of the Doppler signal is approximately 332 m s^{-1} per pixel. Isolating this to include only pixels contained in the linear portion of the absorption line (approximately 700 pixels), our instrument produces a 12.7 m s^{-1} sensitivity per image. Expanding this to include all images in 24 nights of observation produces a full time series with an estimated sensitivity level of approximately 10 cm s^{-1} .

CHAPTER 3

ON AN OBSERVATIONAL CAMPAIGN FROM MT. HALEAKALĀ, HI

This chapter contains material published in Gullledge & Shaw (2022), Frontiers in Astronomy and Space Sciences, Astronomical Instrumentation Section.

To answer the questions which have been posed through use of this newly developed instrumentation, we were awarded six weeks of dedicated observational time on the AEOS 3.6m telescope located at Haleakalā Observatory in Maui, HI in July and August of 2020. During this timeframe, only one night of dedicated observational time was lost to inclement weather, and the observational window was fortunate to have sub-arcsecond seeing throughout the entire duration. These awarded nights were divided between instrumentation installation, alignment, testing, and data collection, and resulted in 24 final, consecutive nights of Jovian observations. During each night, up to six hours of 28-second exposures (with 2 additional readout seconds to maintain timing consistency, resulting in a 30-second cadence) in three simultaneous passbands. Choosing a 28-second exposure time provides sufficient S/N in the darkest (and yet, yielding the highest sensitivity on the disk) region of the absorption line and fine sampling for any existing Jovian oscillations, while also ensuring the exposures remain short enough so that Jovian features do not get blurred as the field rotates. The collected observations provide an average of 200 pixels spanning the diameter of Jupiter, resulting in a theoretical spatial resolution of 0.25 arcseconds per pixel. Utilizing the full 24-day campaign provides a frequency resolution of 0.48 μHz .

3.0.1 Observing Conditions

Detecting the global modes of Jupiter to further our knowledge and understanding of the deep interior was the objective these dedicated observations, and as such, we structured our target list to prioritize collecting data on Jupiter each night. These observations would begin at sunset, and continue every 30 seconds until Jupiter reached 10° above the horizon, at which time we would terminate our exposures due to significant atmospheric effects and horizon obscurations. Because we were awarded entire nights (sunset to sunrise), the remainder of each night was dedicated to observing other planetary targets of opportunity which could help to progress the fields of planetary interior and atmospheric sciences, given the unique capabilities of the PMODE instrumentation.

Ideally, we would hope to collect observations on Saturn during any remaining awarded time, as oscillations have been definitively detected within Saturn’s rings (Fuller (2014), Mankovich & Fuller (2021)), rendering it immensely valuable for further observations with PMODE. However, during the extent of the observational campaign, Jupiter and Saturn were separated by less than fifteen degrees. Therefore, Saturn becomes unobservable less than an hour after Jupiter.

Instead of dedicating less time to multiple targets, we chose to forgo Saturnian observations during the late hours of our awarded nights. We chose to dedicated these hours to collect polarimetric and velocimetric data of Uranus, as it is a significant target of interest in the planetary community, with future proposed National Aeronautics and Space Administration (NASA) flagship missions, and there is still a great deal that is unknown

regarding its atmospheric dynamics and temperatures (Fletcher et al. 2020b). Better constraining these effects will help advance knowledge in an understudied region within planetary science—that of the outermost ice giants (Fletcher et al. 2020a).

The first fourteen of the 45 awarded observational nights were dedicated to the full build and fine alignment of the optical bench for PMODE, as it was necessary to have on-sky light propagating throughout the system to reach the careful alignment necessary. Seven total nights were lost as a result of weather and COVID-19 related closures. The remaining 24 nights were dedicated solely to data collection, spanning 6 August 2020 through 30 August 2020. During this timeframe, we were fortunate to lose only one night to poor weather, 30 August 2020. With the exception of this night, the only breaks in constant observation during this timeframe were due to the unavoidable daytime gaps.

These observations span an Earth-Jupiter velocity range of 11.12 km s^{-1} to 20.37 km s^{-1} . The final planetary dataset consists of (in both Doppler imager channels, and the additional, simultaneous polarimeter channel) 137 total hours dedicated to Jupiter observations, 25 hours dedicated to Uranus observations, and a spare few hours for observations of Mars and Venus once Jupiter, Saturn, and Uranus had each set below the horizon. For this dissertation, we will not focus heavily on the secondary targets for the sake of accurately refining and studying the Jovian dataset. However, it is important to highlight that we have produced a novel and very rich multi-target, multi-band planetary dataset for further study in the astronomical community.

During the Jovian portion of our awarded observations, we experienced an average seeing

value of ~ 0.85 arcseconds atop Mt. Haleakalā. For a seeing of one arcsecond, a spatial resolution of ~ 2800 km per pixel is achieved. For our average seeing of 0.85 arcseconds, we reach a resolution of approximately 2380 km per pixel on the Jovian disk. The final, overall fill factor (or percentage of total time, including daylight gaps) for our collected Jovian data (following elimination of unsatisfactory data with undesirable effects) was 21.79%. Table 3.1 provides a detailed description of the observational conditions solely for the Jovian data, to include timing for observations beginning and ending, along with the total number of nightly exposures, and the number of exposures which were deemed as sufficient quality to remain in our dataset for future analysis. (This was determined to be data counts falling within our predefined cutoffs of 3×10^3 to 4×10^{-4} average counts. These defined cutoffs allow us to eliminate any frames which may be too dim—for example, due to cloud cover within the frame or unexpected dome closures, and frames which may be too bright—for example, due to doors opening into the Coudé room, or computer screens within the room unexpectedly lighting up.)

The table also includes the number of observed hours per night, and the average weather conditions and seeing for that timeframe. Also included is the Jovian apparent magnitude and angular diameter, relative velocity for Jupiter as observed from Earth, phase angle of Jupiter, and our Data Number (DN)-to-velocity conversion unit (our “MOF Unit”). We note that there is a noticeable increase in the value of the MOF unit on the night of 13 Aug 2020 and that the values remain increased following this night. On this date, a fore-optic element in the light path was changed at the site, which manifested as a change in intensity

at the detector, and thus an overall change in the value of the MOF unit.

Starting Date (UT) (year = 2020)	Ending Date (UT) (year = 2020)	# of Acquisitions	# Selected	# of Observed Hours	Average Weather Conditions	Average Seeing (arcsec)	Apparent Visual Mag	Angular Diameter (arcseconds)	Relative Velocity between Jupiter Center and Observer (km/s), positive = away negative = towards	Phase Angle	"MOF Unit" (DN/m/s)
07 Aug 08:50:00	07 Aug 13:33:05	568	567	5	Clear, low humidity, steady wind low, some wind gusts	1.38	-2.701	46.7895	11.1202	4.8609	2.1773e-4
08 Aug 08:46:00	08 Aug 13:15:00	539	539	5	Clear, low humidity, low wind	No Data	-2.697	46.7169	11.5713	5.0491	2.1382e-4
09 Aug 06:29:00	09 Aug 13:08:31	800	792	7	Clear, humidity started low & increased through night, low wind	0.75	-2.692	46.6416	12.0175	5.2354	2.0484e-4
10 Aug 06:24:59	10 Aug 13:29:00	849	841	7	Clear, low humidity, low wind	0.66	-2.688	46.5637	12.4600	5.4195	2.1503e-4
11 Aug 06:39:00	11 Aug 13:20:31	804	745	7	Clear, low humidity, low wind	0.54	-2.683	46.4833	12.8987	5.6015	2.0698e-4
12 Aug 06:12:59	12 Aug 13:04:00	823	814	6.5	Clear, low humidity, low wind	0.52	-2.678	46.4004	13.3307	5.7812	2.0876e-4
13 Aug 06:20:00	13 Aug 13:04:31	810	748	6.5	Clear, low humidity, low wind	0.51	-2.673	46.3151	13.7612	5.9587	2.8408e-4
14 Aug 06:05:00	14-Aug 13:15:00	861	840	7	Clear, low humidity, low wind	0.45	-2.668	46.2275	14.1823	6.1338	2.9457e-4
15 Aug 07:33:59	15 Aug 12:59:30	652	458	4	High humidity, high-altitude clouds	0.53	-2.662	46.1376	14.5998	6.3065	2.9687e-4
16 Aug 06:00:59	16 Aug 12:56:30	832	831	7	Clear, low humidity, low wind	0.62	-2.657	46.0454	15.0104	6.4768	2.8646e-4
17 Aug 07:21:00	17 Aug 12:39:30	638	620	5.5	Clear, avg humidity, high wind gusts	0.56	-2.651	45.9511	15.4168	6.6445	2.9030e-4
18 Aug 05:57:00	18 Aug 12:51:00	829	829	8	Clear, low humidity, low wind	0.81	-2.646	45.8547	15.8175	6.8096	2.7620e-4
19 Aug 05:56:00	19 Aug 12:43:01	815	809	6	Light clouds, avg humidity, low wind	0.38	-2.640	45.7563	16.2103	6.9722	2.7829e-4
20 Aug 06:01:00	20 Aug 08:27:00	293	44	0.5	Light clouds, high humidity, low wind	1.01	-2.634	45.6559	16.5995	7.1320	2.7529e-4
21 Aug 05:47:00	21 Aug 12:21:01	789	357	6.5	Cloudy, high humidity, low wind	No Data	-2.628	45.5536	16.9779	7.2891	2.7229e-4
22 Aug 06:40:00	22 Aug 12:28:31	698	694	7	Light, high clouds avg humidity, low wind	0.84	-2.622	45.4495	17.3507	7.4434	2.5589e-4
23 Aug 06:07:00	23 Aug 12:27:30	762	761	6.5	Clear, low humidity, low wind	0.52	-2.615	45.3436	17.7139	7.5949	2.7108e-4
24 Aug 05:33:59	24 Aug 13:03:30	900	448	4	Clear, low humidity, low wind	0.81	-2.609	45.2361	18.0720	7.7435	2.6637e-4
25 Aug 10:20:00	25 Aug 12:15:00	231	171	2	High wind, high humidity, high altitude particulates from California wildfires	No Data	-2.603	45.1270	18.4221	7.8893	2.7509e-4
26 Aug 06:24:00	26 Aug 12:08:01	689	689	6	High wind, avg humidity, high altitude particulates from California wildfires	1.89	-2.596	45.0163	18.7638	8.0321	2.4744e-4
27 Aug 05:48:00	27 Aug 12:07:31	760	760	6	Clear, low humidity, low wind	2.14	-2.589	44.9042	19.1008	8.1719	2.7567e-4
28 Aug 06:07:59	28 Aug 12:03:30	712	711	6	Some clouds, high wind, high humidity	1.39	-2.583	44.7907	19.4271	8.3088	2.4952e-4
29 Aug 06:21:00	29 Aug 12:06:00	681	670	5.5	Clear, high wind, low humidity	0.62	-2.576	44.6759	19.7480	8.4426	2.4189e-4
31 Aug 06:48:59	31 Aug 12:38:30	700	616	5.5	Clear, low wind low humidity	0.9	-2.562	44.4425	20.3663	8.7013	2.4054e-4

Table 3.1 Jovian observation statistics collected during the Summer 2020 PMODE observational campaign on the AEOS 3.6m telescope. For the sake of this dissertation, low humidity is defined as below 20%, and high humidity is defined as greater than 60%. All statistics pertain directly to Jupiter, as other observed targets are not evaluated to any significant extent within this dissertation.

3.0.2 On-Sky Optimization

A unique feature of the MOFs used for the project is that their respective transmission profiles are variable, and can be tailored to create the desired bandpass by adjusting both the temperature and the magnetic field applied to the vapor cells. The magnetic fields are only adjustable through installation of the permanent rare-earth magnets surrounding the

assembly. For the vapor cells we used during the observational campaign, these magnetic fields were defined to be 2 kilogauss for the 770 nm K MOF cell, and 3 kilogauss for the 589 nm Na MOF cell. The temperature applied to the vapor cells, however, is purely variable by whatever heat (which importantly, must remain stable) that the observers choose to apply during the observational campaign.

A higher temperature causes the single passband to eventually split into two separate passbands. This maximizes the overall possible spatial coverage, but decreases the sharpness in slope of the absorption line, subsequently reducing its sensitivity, as well. Therefore, it is important to fine-tune the temperatures applied to both channels to simultaneously maximize spatial coverage and sensitivity, while also ensuring that the applied temperature is not so high that it will cause the cell to “fail”, or result in the internal vapor depositing on the cell structure itself.

To evaluate the correct temperature to apply to the MOF assemblies, it was necessary to dedicate a portion of our on-sky alignment time to evaluate how varying the temperature results in changing passbands on the disk of Jupiter. We dedicated a portion of our calibration time prior to initial data collection to record images of Jupiter through our complete setup, collecting frames over a range of applied temperatures in small, 1–2°C steps.

Once these data were collected, they were quickly reduced, co-registered, and the Dopplergram (the result following division of the MOF image by the continuum image) was calculated. Within each Dopplergram, we plotted a simple horizontal cut through the center of the Jovian disk to clearly visualize the absorption line profile, we then applied a small

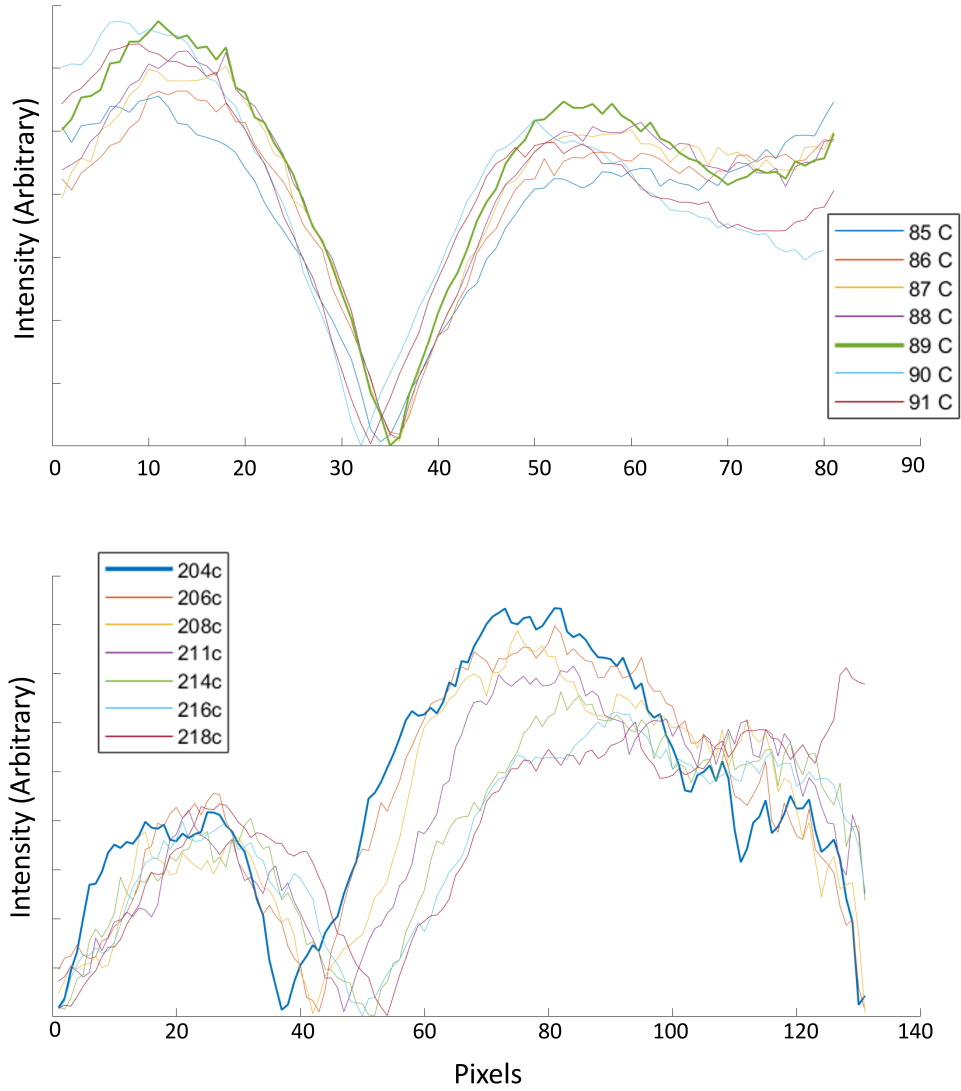


Figure 3.1 Determination process for optimal temperature to apply to the MOF vapor cells. The top panel displays a range of temperatures for the 770 nm K cell, and the bottom panel displays a range of temperatures for the 589 nm Na cell. In both, the absolute value of the derivative of a smoothed cut across a Dopplergram is plotted for each temperature. The line displaying the highest values for the highest percentage of time is determined to be the optimal temperature, and is plotted with a bold line in each panel.

spatial smoothing filter (a moving mean with a width of 30 samples) to the cut. Following the smoothing, we then calculate the derivative of the cut, as a measure of sensitivity (again, as sensitivity scales with the slope of the line), and finally, plot the absolute value of the

derivative across the disk. The cut which displays the highest value of the absolute value of the derivative over the longest region can then be taken as that with the best sensitivity. These cuts are shown for a range of temperatures for both the 770 nm K channel and the 589 nm Na channel in Figure 3.1, where the bolded line in each is the selected final temperature.

For the potassium channel, this process provided a recommended temperature of 89°C applied to the vapor cell, and 204°C for the sodium channel. The single absorption line profiles belonging to each of these two respective temperatures are displayed together within Figure 3.2. The determined profile for the potassium channel has a narrow profile with a steep slope, maximizing the possible sensitivity, but limiting the spatial coverage on the disk.

To compliment this, the temperature chosen for the sodium channel produces an absorption line with a shallower slope, but covering a wider area on the Jovian disk with velocity sensitivity, thus increasing overall spatial coverage on the disk. The combination of both profiles should provide sufficient spatial coverage and high local sensitivity spanning the majority of Jupiter which is observable from Earth (the coverage and sensitivity is representative of Figure 2.6’s sensitivity models).

3.1 PMODE Data Reduction

We chose to develop our data reduction and analysis pipelines in **MATLAB**, as it has a bevy of useful packages and built-in applications to assist with both image registration and signal processing. For this project, all reduction and analysis software required full development from the ground up—as Jovian seismology is a field just beginning to emerge from infancy,




Figure 3.2 Final Determined Temperatures for the K and Na MOFs (as from Figure 3.1), overplotted to display the respective area covered on the Jovian disk. The “K Cut”, or plot through the absorption line profile for the 770 nm potassium Dopplergram at a temperature of 89° C, displays a narrow profile with a steep slope, limiting its spatial coverage on the disk but also exhibiting a high sensitivity. The “Na Cut”, or plot through the absorption line profile for the 589 nm sodium Dopplergram at a temperature of 204° C, was chosen to display a wider profile with a shallower slope, complimenting the 770 nm potassium line by increasing spatial coverage on the disk, but sacrificing a bit of steepness—and thus, sensitivity—by doing so.

there is no inheritance or established techniques to build from. Additionally, as the instrumentation is novel and developed over the course of this work, even the simplest calibration techniques (biases, dark frames, flat fielding routines) all had to be developed as well.

However, simple calibration is not all that must be completed in order to extract the sensitive velocity signals we are in search of. Our pipeline must also encompass field de-rotation to compensate for collecting observations which suffer from field rotation as a result of collecting data from the Coudé room of an Alt-Az telescope, sub-pixel image registration

to ensure no false velocimetric effects are caused by leak-through intensity changes, sensitive edge detection of an extended, space-based target that blurs with seeing variations, image division, and calculation and subtraction of an average frame to remove undesired velocimetric effects from sources besides Jovian oscillations.

The “residual image” which will be produced *via* subtraction of the average frame is multiplied by the desired spherical harmonic, and the total integrated intensity for this multiplied image is then calculated. Repeating this process for every collected image creates our final time-series, which we then convert to the Fourier domain and begin searching for any excess power or signal which could be attributed to Jovian oscillations.

This is simply a brief summary of the steps which must be applied to our data. Each of these steps is explained in detail in the following sections.

3.1.1 Standard Calibrations

Prior to evaluating any of the more complicated data reduction techniques, we first apply standard calibrations (bias-, dark-, “moon flat”—slightly different from your typical dome flats or sky flats, and thus detailed within Section 3.1.2—and bad pixel correction) to the full frame directly from the camera, which consists of both the MOF and continuum images. An example of the raw data collected directly from the detector is shown in Figure 3.3A.

Although the crossed polarizers surrounding the MOF assemblies are intended to completely extinct any continuum image light from “leaking” into the MOF image side of the frame, there is inevitably *some* light being transmitted. To measure this, we kept our vapor cells turned off (e.g., no heat was applied to the cells, so the light in the wings of the

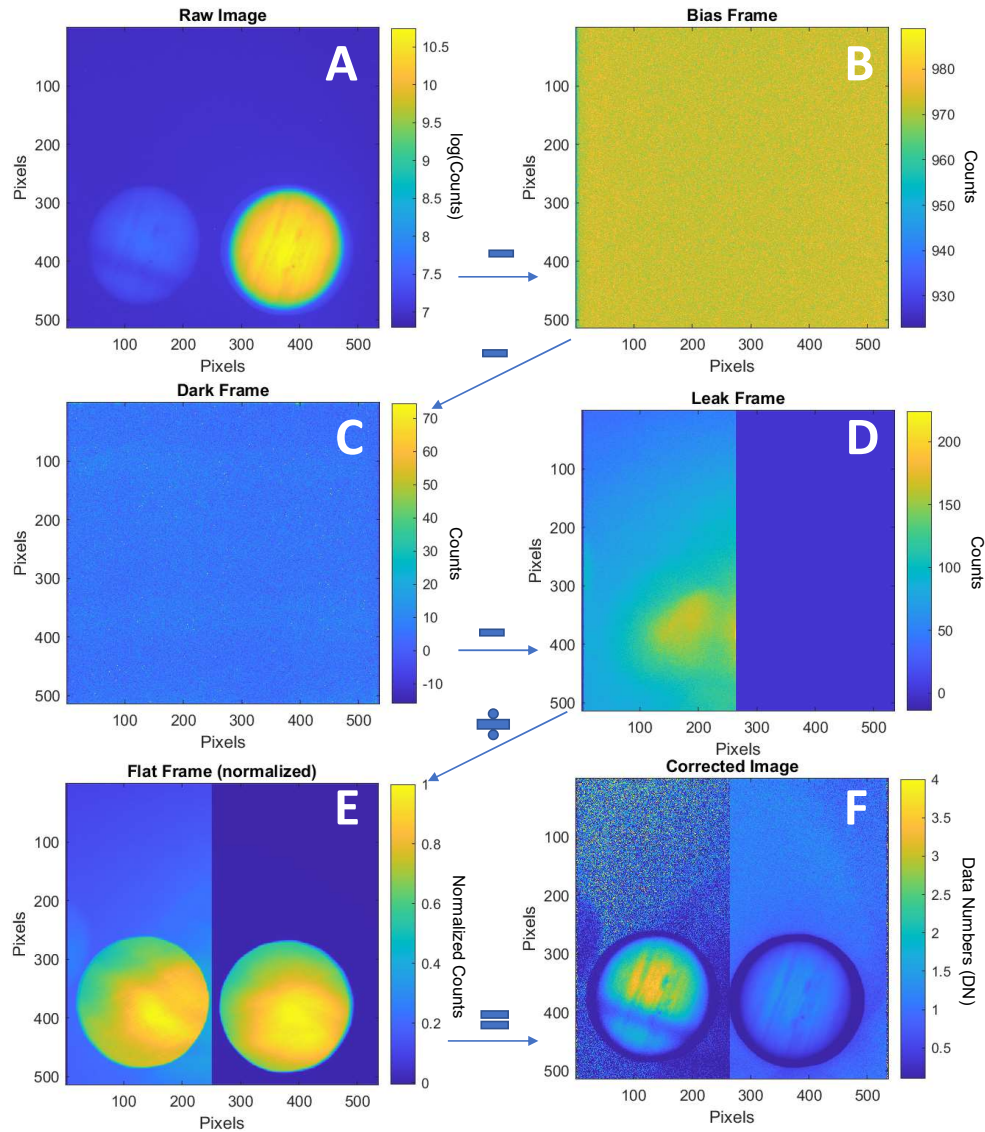


Figure 3.3 Data reduction steps for PMODE observations from the AEOS telescope. Panel A displays a raw image directly off the camera from 06 August 2020, showing how both the MOF image (left side of Panel A) and continuum image (right side of Panel A) are collected on the detector simultaneously. The continuum image is much brighter, so Panel A is displayed with log scaling. Panels B through E show the average bias frame, the average dark frame, the leak image, and the average “Moon Flat” frame (where each side has been normalized to 1 for ease of viewing), respectively. Panel F displays the same image as Panel A, but with linear scaling, after the corrections in Panels B—E have been applied. The MOF and continuum image are now closer in intensity, however the S/N differs.

absorption line should experience no rotation). In this configuration, theoretically no light should be transmitted to the MOF image side of the frame. We then recorded many 28-second exposure frames in this configuration. Averaging these images, and masking out the continuum side of the frame with zeros (therefore no change will be added to the right side of the frame following leak-subtraction), produces a “leak image”, which is shown in Figure 3.3D.

We can see clearly that there is a slight excess of light falling in the same regime as the MOF image. If not correctly removed, this additive noise will not divide out when creating a Dopplergram, thus inhibiting accurate cancellation of Jovian cloud structure. Importantly, the intensity in this “leak image” varies with intensity of the continuum side of the image that is causing the leakage. To compensate for this, we record the average intensity in the *continuum* portion of each frame as basic data reduction steps are being applied.

We divide the average intensity in the leak image by the average intensity in the continuum side of the reduced frame to obtain an array of scaling factors. To avoid any individual image-to-image effects, we then smooth this array of scaling factors with a 30-sample window moving median, and then multiply the leak image by the respective, smoothed, per-image scaling factor. We then subtract this scaled leak image from the data frame, to remove any continuum leakage into the MOF side of the frame.

Following the leak image subtraction, all basic data reduction steps have been applied to our individual frames, which will subsequently be used to produce our final calibrated Dopplergrams. Mathematically, the reduction steps applied to the raw data are as follows:

$$I_{Cal} = \frac{I_{Raw} - B - D - L}{F} \quad (3.1)$$

where I_{Cal} represents the calibrated intensity frame (as in Figure 3.3E), I_{Raw} represents the raw intensity frame (as in Figure 3.3A), B represents the median bias frame (as in Figure 3.3B), D represents the median dark frame (as in Figure 3.3C), L represents the leak image (as in Figure 3.3D) and F represents the median “moon flat” frame (as in Figure 3.3E).

These calibration frames (biases, darks, and leak-images) were collected in bulk during the early nights of the observational campaign, and “quality check calibrations” were obtained periodically throughout the 24 nights to ensure that the calibration frames were consistent and stable. Should any significant variations have appeared within the test calibration frames, the source of the difference would first be determined, then a second set of full calibration frames would be collected. However, these calibrations were stable over the duration of the 24-night observational campaign, allowing our early-run calibrations to be utilized throughout the full dataset.

Following this standard reduction, the image is cropped so that the MOF image (V) and continuum image (C) are isolated, as shown in Figure 2.3. We then divide V by C to produce the Dopplergram (DG), and subsequently divide our Dopplergram by our MOF unit, VF (calculation of this factor is thoroughly detailed in Section 3.3) for final amplitude calibration, as follows:

$$DG = \frac{V}{C} \times \frac{1}{VF} \quad (3.2)$$

3.1.2 Moon Flats

It is important to note that our flat frames differ from the typical “in dome” evenly illuminated flat source (unavailable on the telescope the observations were conducted on), and also differ from typical sky-flats—indeed, the exceptionally narrow pass-band of the MOF technique requires a very brightly illuminated source, and neither dawn- nor twilight-sky frames provided the necessary photons within the MOF image on the detector. Instead, we decide to use portions of the Moon as our flat field. Because the Lunar surface is fairly stable in terms of relative velocity to the Earth, it is useful as a zero point for our further velocimetric analysis.

As the field of view of the AEOS telescope is drastically smaller than the size of the Moon, we were limited to observations of only a small regime of the Lunar surface in any given exposure. The Lunar surface varies significantly in intensity, from dark maria to bright Lunar highlands, each boasting a different albedo. Therefore, we requested that the observers allow the field of view of the telescope to drift over the Lunar surface, while occasionally manually dithering the telescope throughout this drifting. The result of this technique is a dataset of constantly varying images of near-random regions of the Moon, and thus, variations in the Lunar surface average out when we create a “master flat field” from these frames.

Another benefit to our choice of observing the Moon for a period of time with our Doppler imager is the potential to detect reflected Solar oscillations in moonlight, as has historically been done by (Fussell et al. 1995). Naturally, this would provide a strong calibration and

validation source for our instrument. Unfortunately, this goal was not achieved through these observations. As this relies on temporal analysis of individual (or a small number of averaged) images, instead of averaging the entire set, the strong and differing Lunar features in the small field-of-view images we collected proved to inhibit the desired analysis.

In particular, the S/N differences between the MOF and continuum images, and the resultant photon statistics, cause any shadowed regimes in the MOF image (such as lunar maria and dark shadows from craters and ridgelines) to prohibit accurate cancellation of features in the resultant Dopplergram. Thus, the final Lunar Dopplergrams vary in intensity simply due to these S/N variations, which entirely and inevitably mask any true underlying signal that could be due to Solar oscillations.

3.1.3 Field De-Rotation and Image Registration

The MOF image and the continuum image both move within the frame due to tracking and seeing-induced variations throughout the night, although this movement is small due to the inclusion of applied tip/tilt compensation before the light reaches the detector. However, these images move together, maintaining a constant offset between the two images in every frame during the entire observational campaign.

This offset is set by our optical elements, and has been determined empirically to be $(-0.6132, 1.4628)$ pixels in x and y respectively *via* a spiral search through a range of differing shifts, and minimizing the structure in a vertical cut through the Dopplergram following this shift-and-divide process. When the shift is incorrect, more structure will be apparent in this cut; therefore, we minimize the error by finding the shift which most effectively flattens the

cut.

This search provides an offset which we believe is correct to better than 1/20th of a pixel, proof of which is apparent in our full ability to remove effectively all Jovian cloud structure once the continuum image is shifted to the MOF image coordinates and divided out to produce the single Dopplergram (again, an image where the intensity in each pixel pertains to radial velocity towards or away from the observer, which must be multiplied by a scaling factor to translate intensity to velocity) for each frame.

Following calculation of the Dopplergram, we detect the center-of-mass point of each image and shift the image so those coordinates are centralized within the frame as a first rough registration. We then de-rotate each image so that the Jovian North Pole is pointing directly to the center top of our final image frames. This de-rotation is necessary because our observations were collected in the Coudé room of the Alt-Az AEOS 3.6m telescope, and no field-derotation optic is included in the path of light before it reaches the detector.

To calculate the amount each image needed to be rotated, we queried the Jet Propulsion Laboratory (JPL) Horizons ephemerides system (Giorgini 2015) for the duration of our observational campaign, to include the timing, altitude (Alt), azimuth (Az), and declination (δ) of Jupiter, along with the position angle (PA) and the hour angle (θ_H). This information, when matched with the timing information for each of our images and combined with the latitude of our observational site (ϕ), allows for calculation of the parallactic angle for each ephemeris query point as follows:

$$\theta_p = \text{atan2}(\sin(\theta_H \times 15^\circ), (\tan(\phi) \times \cos(\delta)) - (\sin(\delta) \times \cos(\theta_H * 15^\circ))) \quad (3.3)$$

Combination of the Jovian position angle from the ephemeris with this parallactic angle, the Jovian azimuth and altitude, and an empirical constant C (which sets the rotation to the desired area—in this case, North—of the final image frame) provides the amount by which we need to rotate the image in our frame:

$$\theta_{RN} = |\theta_p - az + alt - PA + C| \quad (3.4)$$

Here, θ_{RN} (for Rotate North), represents this final rotation angle. This θ_{RN} is pre-calculated for our entire ephemeris array so that when the de-rotation step is reached in the reduction pipeline, the array can simply be queried to determine which time in the array is closest to the true observational time, and the respective θ_{RN} value can be extracted. Importantly, to avoid any interpolation or edge effects, we upscale each 241×241 pixel Dopplergram by a factor of four, then rotate it by the desired amount with a nearest-neighbor interpolation, before then scaling our image down by a factor of four, back to its original size.

Following this de-rotation, each image of Jupiter should now be in the rough (within 1–2 pixels) of the center of the frame, with its major axis spanning from East to West in the frame, and its minor axis spanning from North to South. This ensures that the subsequent frame-to-frame registration will be matching similar images without having to take rotation

effects into consideration, a useful simplification to the registration process.

The frame-to-frame registration throughout the night is vitally important to avoid noise in our time series, so we apply a multi-step registration process. First, we use a zero-crossing Canny edge detection algorithm intrinsic to the **MATLAB** image registration toolbox. We chose to implement this zero-crossing edge detection technique because these techniques are insensitive to seeing-induced blurring, and have historically been successfully implemented in helioseismology (Toner & Jefferies (1993), Hill et al. (1975)). The Canny algorithm, in particular, was chosen for producing results similar to the classic Laplacian of a Gaussian registration technique, but with enhanced detection and localization performance (Canny 1986).

Once this edge has been detected, we mask out the background surrounding the edge so that all that remains within the frame is the Dopplergram. Following this masking, we apply a monomodal intensity-based registration algorithm, also intrinsic to the **MATLAB** image registration toolbox, to register each Dopplergram to the exact center of the frame.

The monomodal option was chosen because it is designed to work for images with similar intensity and contrast that are collected on the same detector, and subsequently, well-suited for our registration attempts. Overall, the frame-to-frame registration variability produced is smaller than 0.04 pixels; we believe this to be a sufficient stability, and thus, move forward with all Jovian images are co-located for future average subtraction and temporal analysis.

All above data reduction steps are applied in the pipeline included in Appendix A.

3.2 PMODE Residual Generation

Now that our data has been accurately reduced, it is important to ensure that any external velocity signals are removed from our data. This is exceptionally important, because if we are looking for signals at or below the level of $\sim 50 \text{ cm s}^{-1}$ (Gaulme et al. 2011), but other velocimetric signals are superimposed—such as the Jovian rotation, with a velocimetric difference of 25.4 km s^{-1} (four orders of magnitude higher than that of the reported oscillations oscillations!)—these will inevitably dwarf any signals due to the global modes, if not properly removed.

3.2.1 *Additional Velocity Calibrations*

We know the measured Doppler shift of any reflected Solar line will intrinsically be a combination of many different velocimetric sources. First, and most significantly varying, is the velocimetric change caused by Jupiter’s rapid rotation, $V_{J,rot}$. As described in Section 2.2.2, this effect is what restricts the sensitive region to only the narrow slice of the absorption line on the Jovian disk, as the remainder of the disk is Doppler shifted out of the width of the absorption line. This is the only spatially varying velocity effect which must be removed, and appears in our data as a gradient in Doppler shift, which spans a range from -12.57 km s^{-1} to $+12.57 \text{ km s}^{-1}$ over the Jovian disk.

Secondly, the next clearly changing velocimetric effect is caused by the relative motion in the Earth-Jupiter orbital system, hereafter $V_{J/E}$, which appears as a change in the central location of the absorption line profile on the Jovian disk. At opposition, this absorption

line will fall directly in the center of the Jovian disk. As we move away from opposition, and Earth effectively speeds away from Jupiter, the relative velocity between the two bodies increases, and the absorption line center shifts towards the edge of the Jovian disk.

Over the duration of our 24-night observational campaign, the Earth-Jupiter relative velocity increased from a minimum of 11.12 km s^{-1} to a maximum of 20.37 km s^{-1} , as obtained from JPL Horizons (Giorgini 2015). This 9.25 km s^{-1} velocity change causes our absorption line center to shift from slightly left-of-center on the Jovian disk on the night of 06 August 2020, to nearly the edge of the Jovian disk at the end of the observational campaign, on 30 August 2020.

Figure 3.4 displays this effect with a representative MOF image pulled from our bias-, dark-, flat-, and leak-corrected dataset. These images also display the slight change in the extent of Jupiter itself on the detector (beginning with a diameter of 204 pixels at the start of the observational campaign and decreasing to a value of 196 pixels at the end of the campaign), which also occurs as a result of the increased distance between the Earth and Jupiter as the relative velocity increases. This $V_{J/E}$ factor is uniform over the Jovian disk and is on the order of \sim a few km s^{-1} . Although the value is increasing as the observational campaign progresses, it is fairly constant throughout one single night, and will be treated as such for the remainder of our analysis.

Another velocimetric effect in the Earth-Jupiter orbital system is that of the Earth's nightly rotation, $V_{E,rot}$, which appears as a predictable and gradual variation in intensity throughout the night, on the order of hundreds of m s^{-1} ,




Figure 3.4 Change in the center location of the 770 nm potassium absorption line center, induced from the Earth-Jupiter system relative velocity changes over the course of the 24-night PMODE observational campaign during the summer of 2020. These three images display a MOF image pulled from our dataset for the first night of the campaign (06 August 2020), the central night of the campaign (18 August 2020), and the final night of the campaign (30 August 2020), along with the corresponding relative velocity between the Earth and Jupiter for each respective night. The dark band of the absorption line can clearly be seen shifting from the left third of the planet to the far limb of the planet as time progresses, thereby changing the spatial coverage on the disk contained within the velocity-sensitive absorption line. All images have basic calibrations applied, and each is normalized to 1 for comparison. Each image has approximately 200 pixels spanning the equator (specifically, ranging from 204 pixels at the beginning of the campaign to 195 pixels at the end of the campaign.)

Similarly to the Earth-Jupiter relative distance effect, we also must consider the velocimetric effect which results from the Jupiter-Sun orbital system (hereafter $V_{J/\odot}$). This manifests exactly as $V_{J/E}$ does, as a factor uniform across the disk and appearing again as a slight offset in the position of the absorption line; however, the effect is much smaller (a few tens of cm s^{-1}) because Jupiter is orbiting much more slowly than the Earth.

Finally, velocimetric effects intrinsic to Jupiter will be included in the Doppler shift of the

absorption line as well, including winds (both zonal and meridional, represented as $V_{J,Wind}$ which should be hundreds of m s^{-1}), upwelling, storms, and finally, any global oscillations themselves, represented as V_{Osc} , theoretically with a maximum of $\sim 50 \text{ cm s}^{-1}$, per previous reports (Gaulme et al. 2011).

Combined, all of the effects influencing the Doppler shift of the absorption line is:

$$V_{Doppler} = V_{J/\odot} + V_{E,rot} + V_{J/E} + (1 + \cos(\phi)) (V_{J,rot} + V_{J,Wind} + V_{Osc}) \quad (3.5)$$

The factor of $(1 + \cos(\phi))$ which multiplies all factors that are intrinsic to Jupiter (such as rotation, winds, and oscillations) compensates for a doubling effect in Doppler velocity caused by reflection of light off of the Jovian cloud decks. Here, ϕ represents the phase angle between Jupiter and the Earth, which is approximately 2 for the time period that our observations were collected. A more thorough discussion of these additive velocity effects can be found in Gaulme et al. (2008), Gonçalves et al. (2019), and Cacciani et al. (2001).

Importantly, all of these effects (with the exception of the Earth's rotation and the intrinsic oscillations) change over long timescales. Because of this, they can be excluded from our data by calculating an average Dopplergram for each night (as shown in Figure 3.5B, which will contain $V_{J/\odot}$, $V_{J/E}$, $V_{J,rot}$, and $V_{J,Wind}$. This average frame can then be simply subtracted from each individual Dopplergram to remove the slowly-changing additive velocity effects.

We expect this subtraction of the average Dopplergram to successfully eliminate the majority of the additional velocity sources, while leaving any true, short-term oscillatory signals unaffected. If we are interested in examining the Jovian atmospheric dynamics instead, we

will search for their evidence within each nightly average, as subtraction of these will also remove signals from upwelling and zonal winds from the residual Dopplergrams.

The resultant frame after subtraction of the Dopplergram is our “residual image”, an example of which is shown in Figure 3.5C. Importantly, in this figure, the strong gradient of the absorption line profile from panels A and B has been removed, producing a residual which appears more as static noise. However, we can now see that a change in sign is induced from the absorption line, as the left half of the absorption line region is lighter in the residual image, and the right half is darker.

This is because the gradient of the absorption line is negative on the left side, and positive on the right side. This effect is removed through later division by our “MOF Unit” image, which effectively multiplies the left portion of the absorption line region by a value of -1 to compensate for this effect and ensure that intensities in both sides will vary cohesively, as opposed to a net cancellation effect.

Prior to calculation of the *final* average Dopplergram, we wish to remove the lower-quality data which is significantly degraded by atmospheric effects at the end of each night, which decreases in intensity as a result of the light propagating through a higher airmass as Jupiter approaches the horizon. To do so, we calculate a *full* preliminary average Dopplergram encompassing all frames, including those in the late night, and subtract this preliminary average from all individual frames.

We then integrate the remaining intensity in these preliminary residuals, shown in Figure 3.6. This plot clearly shows that the ~ 70 frames at the end of the night differ from the aver-




Figure 3.5 Panel A: An individual Dopplergram, as in Figure 2.3C. Panel B: The average Dopplergram for a single night, calculated as the median of all individual Dopplergrams. Panel C: A residual image, calculated by subtracting the average Dopplergram (Panel B) from the individual Dopplergram (Panel A). All three images are in “Data Number” (DN) units, prior to conversion to velocimetric units.

age fairly significantly, displaying the falloff due to seeing-induced variations. From empirical testing, we find that this falloff for each night can be isolated by smoothing this restricted set of preliminary residuals, and then calculating the gradient for this full, smoothed set, and defining the cutoff as the location where at least the subsequent ten frames exceed two standard deviations from the median value of the residual gradient. An example displaying the location of a cutoff index for one night is shown in Figure 3.6, where the index is represented as a black line. Following removal of late-night data, our final average Dopplergram can now be calculated as the median of all remaining frames.

3.2.2 Dopplergram Cleaning

Once the final average Dopplergram has been calculated we wish to remove any insensitive portions of the Jovian disk, as these will simply add noise if included in our final calculation.




Figure 3.6 Representative night of preliminary residuals to display an example of late-night data which will be removed for further calibrations. This dataset shows integrated intensity in a set of residual images obtained by calculating a single, full nightly average, subtracting this average from every frame, and integrating the resultant values. We then smooth this set, calculate the derivative, and define the cutoff index (represented as the black line) as where our data begins to regularly exceed more than two standard deviations from the median value of the residual gradient. All data following this index (represented here in orange) is removed and not considered in any further calculation.

To do so, we average the central twenty horizontal cuts through our average Dopplergram to produce an average cut across the disk, which is shown projected on the disk in the leftmost panel of Figure 3.7.

We define our cutoff to be 80% of the maximum depth of the absorption line, when compared to the median value of the insensitive side of the disk. This removes the region outside the linear portions of the absorption line (these linear portions are where our sensitivity is contained). The central panel of Figure 3.7 displays the projected cut, with added lines showing the minimum of the cut, the median value of the insensitive region of the Jovian disk, and our defined 80% cutoff. Once this cutoff has been determined, we next remove all values which are higher. Finally, we also mask the central ten pixels of the absorption line, as the slope is effectively zero here, and thus, the sensitivity is low as well. The average




Figure 3.7 Left Panel: The nightly averaged Dopplergram, with the average central cut overlaid along the disk. Central Panel: The average cut, with the minimum, median, and 80% cutoffs represented, to show isolation to only the portions of the absorption line profile which exhibit a linear slope. Right Panel: The same average Dopplergram as in the rightmost panel, following removal of the outer 10% of the Jovian diameter and isolation to the linear portion of the absorption line. The color scaling for the left and right plots, and the y-axis for the central plot, are all represented in units of Data Numbers, prior to MOF unit conversion into velocimetric units.

Dopplergram, after restriction to these defined sensitive areas, is shown in the rightmost panel of Figure 3.7.

The indices of this cleaned average frame are recorded, and all pixels in each individual frame not falling within this isolated region are masked out accordingly. Finally, after each individual Dopplergram has been cleaned, the average Dopplergram is subtracted from each frame to produce the individual residual images. These must now be calibrated so that their values correspond to Doppler velocities, as opposed to simple data numbers.

All above residual generation steps are applied in the pipeline included in Appendix B.

3.3 Calibration of Raw MOF Data to Velocity Units

Section 2.2.2 has already discussed some of the effects which occur as a result of Jupiter's rapid rotation, which results in a Doppler velocity signal of $\pm 12.6 \text{ km s}^{-1}$ on each side of Jupiter at the equator. However, when viewing reflected light from the Sun-Jupiter-Earth system, a doubling effect (as earlier detailed in Section 3.2.1) occurs, creating a velocimetric change spanning the range of $\pm 25.2 \text{ km s}^{-1}$ across the sunlit side of the Jovian disk. The knowledge of this value allows us to calculate necessary velocity calibration values for our instrument through use of in-image sources.

3.3.1 Original Calibration of Raw MOF Data to Velocity Units

In particular, we know that our output images contain ~ 200 pixels across the Jovian disk—these 200 pixels correspond directly to the $\pm 25.2 \text{ km s}^{-1}$ induced from Jupiter's rotation from edge to edge. This spatial extent on the detector, combined with the rotation rate, can be used to provide a value in $\frac{\text{m/s}}{\text{pix}}$. We then calculate the intensity change within the absorption line profile of the Dopplergram (following bias-, dark-, leak-, and flat-correction, and division of the MOF image by the continuum image) corresponding to a range of pixels (we chose 15 for consistency), to obtain a value of $\frac{DN}{\text{pix}}$, with DN representing the values in the Dopplergram following all calibrations, corrections, and image division.

Combining these two values provides a final calibration unit of $\frac{DN}{\text{m/s}}$. This calibration can then be divided into each pixel value to convert from DN to m s^{-1} . The final equation for

calculating this unit for each night is as follows:

$$VF = \frac{\Delta_f}{R_{size}} \times \frac{Jup_{size}}{4V_{Jup}}, \quad (3.6)$$

Here, VF represents our MOF unit in $\frac{DN}{m/s}$, with Δ_f representing the intensity variation along the measurement region R on the Dopplergram with size R_{size} (in pixels), Jup_{size} being the Jupiter diameter in pixels on the detector, V_{Jup} being the rotation velocity on Jupiter at the equator (in $m\ s^{-1}$), and the (approximated) factor of 4 comes from the velocity difference at each edge of Jupiter equal to $+V_{Jup}$ and $-V_{Jup}$, multiplied by the phase factor $(1 + \cos\phi)$, which again, is almost equal to two for observations from the Earth near opposition.

This MOF unit varies slightly from night to night, primarily due to the shifting of the absorption line center as a function of the radial velocity between Jupiter and Earth (an effect visualized in Figure 3.4, and discussed further in Cacciani et al. (2001)), in addition to smaller second-order effects (such as oblateness, differential rotation, and inclination, which are not considered in this current analysis but remain for future investigation), and is calculated accordingly for each night to account for this change.

With this single unit calibration, it is subsequently necessary to multiply all pixels on the leftmost half of the absorption line by -1 as compensation for an induced change in sign when crossing the bottom of the absorption line, because the gradient on the left side of the line is negative, while the gradient on the right side is positive, and it is necessary to ensure that both sides of the line are changing in tandem, instead of in opposition.

3.3.2 Improved Calibration of Raw MOF Data to Velocity Units

This calibration technique produces one single value, which is true over the 15 pixel range for calculation—however, it is more accurate to calculate this value for every single pixel within the image. To do so, the exact same procedure as detailed above is applied—however, instead of approximating the gradient over a 15 pixel regime for a few averaged row across the average velocity image (Figure 3.8A), we now isolate the average image to only the absorption line itself, where the rest of the disk has been discarded (Figure 3.8B). For each row in the image, we take a single cut (Figure 3.8C), then obtain the gradient for that particular line. This gradient itself now has values of $\frac{DN}{pix}$; therefore, is then multiplied by the $\frac{Jup_{size}}{4V_{Jup}}$ factor, so that the gradient line is now in the “MOF Unit” values, or VF , which again, represents our MOF unit in $\frac{DN}{m/s}$ (Figure 3.8D).

This process is repeated for every single row in the cleaned average image, until a “gradient MOF Unit image” is produced (Figure 3.8E). This calibration image can then be divided into each true science image for its particular night, calibrating the image values (initially in DN) to velocity values ($m\ s^{-1}$). A representation is shown by calibrating the original average velocigram image (Figure 3.8B) by the newly calculated “Gradient MOF Unit Image” (Figure 3.8E), to produce an example calibrated image (Figure 3.8F).

This MOF unit image is calculated *via* this described method for each night, and is then used to calibrate each corresponding science image for that respective night. This new correction method produces a method in which every pixel is accurately converted to velocimetric units, instead of simply dividing each pixel value by the “approximate” MOF

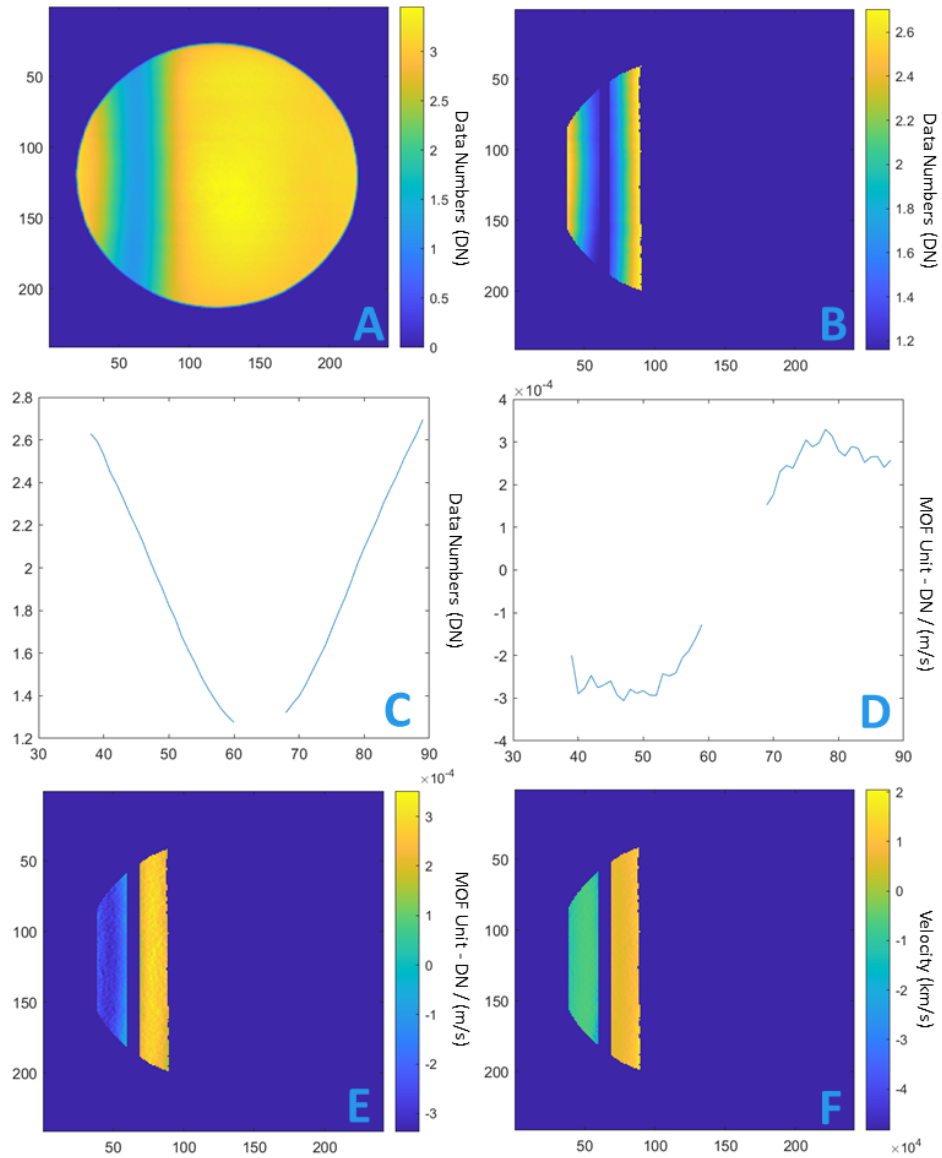


Figure 3.8 Modified technique for calculating the MOF Unit. Panel A: full average velocigram. Panel B: cleaned version of full average velocigram (Panel A). Panel C: Cut through a single row of cleaned velocigram (Panel B). Panel D: Gradient of the single cut through the velocigram (Panel C). E: Full MOF image, from gradients calculated for each row of the average velocigram (Panel B). F: Average velocigram (from Panel B) calibrated to velocity units by the MOF Image (from Panel E).

unit per night as originally explained. Although this modification to the calibration produces a more accurate result, the initial approximation was very close, and the difference between these two techniques is effectively negligible. However, we choose to move forward with the full image correction as opposed to the average region correction, to ensure that our data are corrected as accurately as possible.

With this full image MOF unit calibration, it is no longer subsequently necessary to multiply the leftmost pixels by -1 to compensate sign change at absorption line minima. Now, the “MOF Unit Image” itself remains negative on the left side of the line, and positive on the right side of the line, as shown in Figure 3.8. Therefore, the simple division of each true image by its corresponding nightly MOF Unit Image compensates for the change in gradient direction of the absorption line, while also converting our residual images into accurate velocimetric units. Once this calibration has been completed, we are now ready to multiply each image by our desired spherical harmonic for further analysis.

This above MOF Unit conversion step is applied in the pipeline included in Appendix B.

3.4 Spherical Harmonic Multiplication

Following all aforementioned data reduction steps, we next multiply these residual Dopplergram frames by the spherical harmonic of interest, generated and fully normalized as follows:

$$Y_\ell^m(\theta, \phi) = \frac{1}{\sqrt{2\pi}} P_\ell^m(\cos \theta) e^{im\phi} \quad (3.7)$$

where $P_\ell^m(\cos \theta)$ represents the corresponding Legendre polynomial for a given ℓ and m , and the factor of $\frac{1}{\sqrt{2\pi}}$ is the normalization coefficient, chosen to ensure that the generated spherical harmonics are orthonormal over their full area (before restriction to the isolated regime on the Jovian disk), thereby preserving true amplitudes present within the dataset.

Figure 3.9 displays representations of the projected spherical harmonics, for degree ℓ and azimuthal order m of a set of spherical harmonics, up to and including the degree $\ell=5$, for reference, calculated *via* the Python software *starry* (Luger et al. 2019).




Figure 3.9 Representations of the projected spherical harmonics, for degree ℓ and azimuthal order m of a set of spherical harmonics, up to and including the degree $\ell=5$. Hemispherically antisymmetric spherical harmonics are represented with a dashed circle surrounding them. This figure was generated through implementation of the *starry* software for Python (Luger et al. 2019).

Finally, the intensity in each pixel of the residual velocigram (following multiplication by the spherical harmonic of interest) is integrated to produce a time series of Doppler data

throughout each night. A description of this process for one single pixel is as follows:

$$R = (DG - DG_{Avg}) \times Y_\ell^m(\theta, \phi) \quad (3.8)$$

where R represents residual intensity, DG represents a single Dopplergram, DG_{Avg} represents the average Dopplergram, and $Y_\ell^m(\theta, \phi)$ represents our desired, normalized spherical harmonic. This process is applied to each pixel in the image, and the resultant average of these final pixels is obtained as our final data product to obtain one value, per Y_ℓ^m mode, per image.

All above spherical harmonic multiplication steps are applied in the pipeline included in Appendix C.

3.5 Lomb-Scargle Periodogram Calculation to Obtain Final Amplitude Spectra

Once our time series—consisting of reduced, cleaned, and calibrated integrated intensity for the Y_ℓ^m mode of interest—has been produced by obtaining the single value of interest from each science image, the next step is calculating the Fourier transform of this time series to reveal any periodic signals present within the dataset.

Before converting our data to the Fourier domain, a few final data-cleaning steps must be applied. First, we apply a high-pass robust quadratic regression smoothing filter with a 200-sample window, resulting in a cutoff of 166 μHz to each individual night within our time series (thus, we remove our capability to examine signals lower than 166 μHz).

This smoothing significantly reduces any signals in our data due to periodic upwelling events that we know to be present within the belts and zones, which are expected to con-

taminate our data up to a maximum of $700\mu\text{Hz}$ (Lederer et al. 1995), and also reduces any low-frequency instrumental noise effects, as are expected to be present due to seeing-induced changes throughout the night, in addition to tracking and image registration variations. The goal of this smoothing is not to fully eliminate these signals (both the true velocity ones due to upwelling, and also the instrumental noise effects), but instead, to simply minimize ringing and noise elsewhere in the spectrum as a result of high-power, low-frequency noise.




Figure 3.10 Representation of the Tukey window that is applied to each night of residuals to smoothly roll the values off to zero and avoid sharp transitions at the beginnings and ends of each night. This window has a cosine fraction of 0.1 and calculated at equal length to each night of residuals. Only the ~ 50 values at the beginning and end of each night have reduced intensity, the majority of data points are entirely unaffected.

Once any long-term nightly trends have been removed by this smoothing, we then multiply each night of data by an equal-length Tukey window with a cosine fraction of 0.1 (shown in Figure 3.10) to smoothly roll off the edges of the residuals to zero, creating smooth transitions between our data and observational gaps. Creating these smooth rollofs in the time domain will later minimize edge effects in the Fourier domain. The residuals at each of

these steps (“raw” but reduced and multiplied by the desired spherical harmonic, smoothed, and multiplied by the Tukey window) are shown in Figure 3.11, for the example night of 15 August 2020 for the $Y_\ell^m = (1, 0)$ mode.

Following these pre-processing steps for each individual night of data, the nights are all combined together with zero-values inserted accurately during all daytime observational gaps through use of the exact Modified Julian Date (MJD) values for each night to preserve the true spacing in our dataset. Once all twenty-four nights have been combined into one long time-series string, we then apply outlier rejection to this full string. For this project, we define outliers as values which live more than three scaled Median Absolute Deviations (MAD) from the median value, when considering the entire 24-night observational campaign. Per **MATLAB** documentation, the scaled MAD is defined as:

$$MAD = c \times median(|A_i - median(A)|) \quad (3.9)$$

for $i = 1, 2, \dots, N$, and where:

$$c = \frac{-1}{\sqrt{2} \times \text{erfcinv}\left(\frac{3}{2}\right)}. \quad (3.10)$$

Here, **erfcinv** represents the inverse complimentary error function, and is effectively a simple weighting function to help improve outlier rejection. Figure 3.12 displays the full 24-night campaign prior-to and after the outlier-rejection, which is fairly conservative and preserves the majority of the true variation within the dataset.




Figure 3.11 Preparing individual nights of residuals prior to combination and conversion to the Fourier domain. The top panel displays a set of processed integrated intensity for the $Y_\ell^m = (1,0)$ mode for 15 August 2020. The middle panel displays the residuals after the robust quadratic regression smoothing filter with a 200-sample window has been applied to the top panel to remove overall trends. The bottom displays the residuals from the middle panel after a Tukey window rolloff with a cosine fraction of 0.1 has been applied to smoothly roll the edges of the residuals off to zero, to avoid edge effects and aliasing in the Fourier domain.

Finally, once all smoothing, gap-preserving and combining, and outlier rejection has been applied, we can then convert these data to the Fourier domain, for each Y_ℓ^m mode of interest.




Figure 3.12 Outlier rejection visualization for the full PMODE observational campaign, shown on the $Y_\ell^m = (1, 0)$ mode as an example. Gaps from daytime breaks in the observational campaign are preserved. The blue data is representative of the integrated intensity residuals that have had linear regression smoothing applied, and the orange data is after the 4-MAD outlier rejection has been applied. A zoom-in of one representative night is shown as an inset in the upper left of the plot.

For this conversion, we choose to implement a Lomb-Scargle periodogram (after Lomb (1976) and Scargle (1982)), one of the most frequently and successfully implemented techniques to

detect periodicity in astronomical data, as opposed to a classical Fourier transform. The primary difference is simply that the Lomb-Scargle periodogram can internally handle signals with uneven sampling or with missing samples, something that the strict Fourier transform will struggle with—a detailed explanation of this technique and its value can be found in VanderPlas (2018).

Although our data was set to collect exactly every 30 seconds (with a 28-second exposure and two additional seconds to allow for readout), and started on the exact minute per the NIST Official U.S. time website, we digress that there may be inevitable slippage of up to a couple seconds per night during the campaign. As the exact true time is recorded for each exposure, and this true time is also used as an input for our Lomb-Scargle periodogram, we can now accurately account for any slippage whilst also preserving true signals.

We have carefully calibrated the amplitude of our time series, as detailed in Section 3.3.2. Now, we want to ensure that these accurate amplitudes are preserved and maintained throughout our Fourier analysis techniques. For this, we include a conversion factor to retrieve amplitude from our Lomb-Scargle periodogram as follows:

$$A = \sqrt{\frac{2 \times P \times f_s}{N_T}} \quad (3.11)$$

where A represents calibrated amplitude, P represents the resultant power from the Lomb-Scargle periodogram (which we choose to oversample by a factor of 3), f_s represents the frequency sampling rate (1/30 Hz), and N_T represents the number of samples in the true input time series.

This produces our final amplitude spectra which will be analyzed in detail throughout the

remainder of this manuscript, with amplitude accurately preserved. Finally, each amplitude spectrum is calculated on the range of 0–6000 μHz , as this encompasses beyond the highest frequencies at which models predict global modes can be found (Markham & Stevenson 2018).

All amplitude spectrum generation steps are applied in the pipeline included in Appendix D.

3.6 Analysis of Spherical Harmonics

3.6.1 *Spherical Harmonic Leakage*

As discussed prior, an additional noise complication which must be considered is that of spatial aliasing which occurs as a result of restricting our instrumental sensitivity to only a small band across the disk, as in the case of the 770 nm potassium line. As this is the only dataset we have thoroughly analyzed, due to the undesired detector noise in the 589 nm sodium channel, it is important to examine more closely.

In addition to the spatial leakage (a reminder: measurement of an undesired mode when examining another specific, desired mode), another complicating factor here is that our region of sensitivity on the Jovian disk is not constant throughout the observational campaign. In fact, the location of the line-center shifts each night as a result of the relative orbital velocity between Earth and Jupiter, as detailed in Section 3.2.1.

Over the duration of our 24-night observational campaign, the Earth-Jupiter relative velocity increased from a minimum of 11.12 km s^{-1} to a maximum of 20.37 km s^{-1} , as

obtained from JPL Horizons (Giorgini 2015). This 9.25 km s^{-1} velocity change causes our absorption line center to shift from slightly left-of-center on the Jovian disk on the night of 06 August 2020, to nearly the edge of the Jovian disk at the end of the observational campaign, on 30 August 2020. Figure 3.4 shows that not only is the central location of the line shifting, but the number of pixels contained within the sensitive regime of the absorption line is also decreasing as the line approaches the limb of Jupiter.

The resultant spatial leakage is unavoidable (with the exception of adding a secondary MOF channel, as initially planned). However, we can bias our analysis to those modes which suffer less of an effect from this leakage. To determine which modes we should focus our efforts on, we generate all normalized spherical harmonics of interest, and then—for every night of interest—isolate the spherical harmonic to include only the region of the disk where we have velocity sensitivity surrounding the absorption line; these regions are displayed in the left column of Figure 3.13.

Once the spherical harmonics have been isolated to the defined region of the disk, we can calculate the spherical inner product of that region-isolated spherical harmonic with itself, to determine how significantly restriction to that particular region of the disk is affecting it. Because spherical harmonics are orthonormal *over the entire sphere*, we know that the below equation is always true:

$$\int_{-1}^1 Y_{\ell}^m(x) Y_{\ell'}^m(x) dx = 0 \quad (3.12)$$

when $\ell \neq \ell'$, and

$$\int_{-1}^1 Y_{\ell}^m(x) Y_{\ell'}^m(x) dx = 1 \quad (3.13)$$

when $\ell = \ell'$. When we are restricted to an isolated regime on the sphere, we will instead get a value somewhere between 0 and 1, where the closer said value is to unity, the less affected it is by spatial aliasing. We repeat this process for each desired spherical harmonic: restricting it to the desired region, multiplying the restricted spherical harmonic by itself, and integrating the result.

Leakage matrices calculated in this manner are displayed in the right column of Figure 3.13, for a representative half-disk as a reference (region of interest shown in Panel A, leakage matrix shown in Panel B), then the isolated sensitivity for the first night of Jovian observations (our spatial maximum, region of interest shown in Panel C, leakage matrix shown in Panel D), and then for the final night of Jovian observations (our spatial minimum, region of interest shown in Panel E, leakage matrix shown in Panel F).

In each leakage matrix, the size of the dot corresponds to the percentage of the recovered spherical harmonic, where the smallest dots have correlation values of 0.0161, and the largest dots have correlation values of 0.3301. Larger dots suffer from less leakage and are thus more desirable.

3.6.2 Determination of Final Set of Modes for Analysis

We can see from Figures 3.13D and 3.13F that, as expected, as we lose surface area for our sensitivity, we suffer from more leakage (as evidenced by the overall smaller dots in Panel

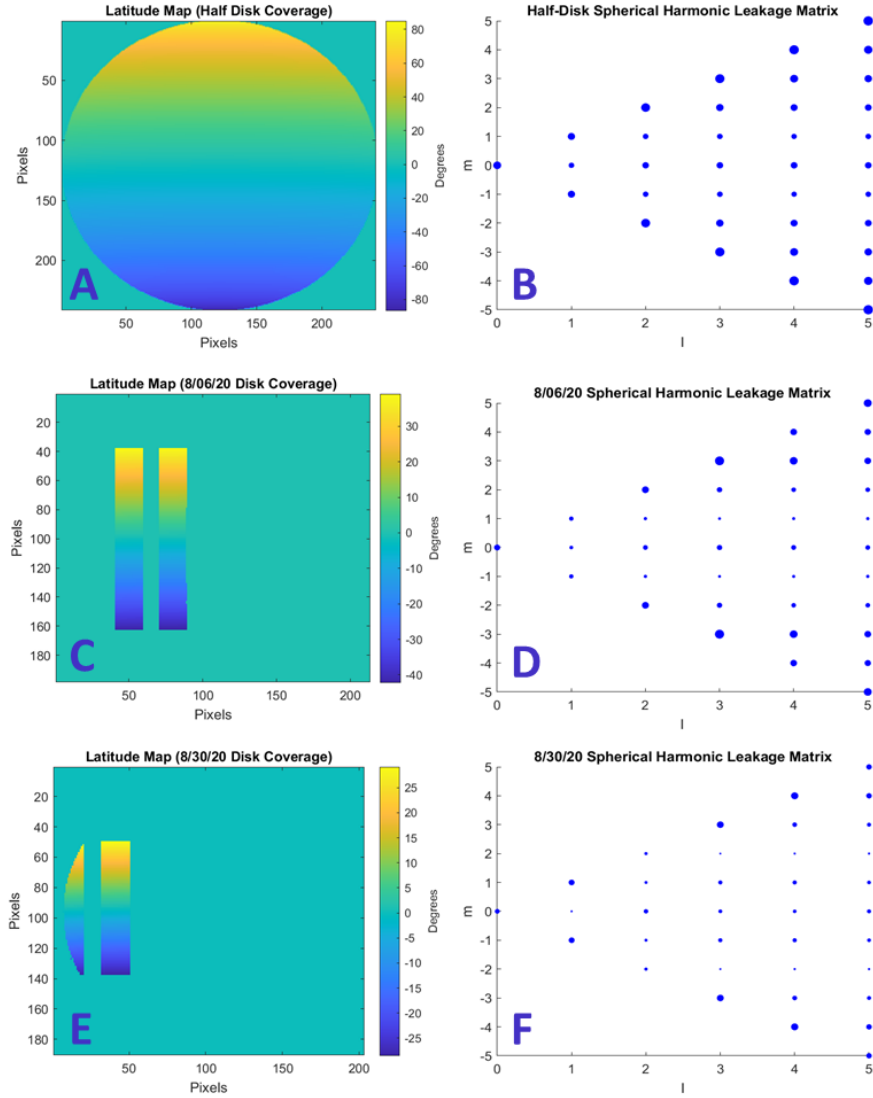


Figure 3.13 Spherical harmonic leakage matrices for the PMODE observational campaign. The left column displays the regions of the disk included for calculation of the leakage matrices, beginning with an entire half-disk in Panel A for reference. Panel C displays the spatial coverage of the 770 nm K absorption line on the first night of observations, and Panel E displays the spatial coverage on the final night of observations. The corresponding leakage matrices obtained by calculating the spherical inner product of each regionally-isolated spherical harmonic with itself are shown in the right column, where the size of the dot corresponds to the percentage of the recovered spherical harmonic, where the smallest dots (Panel F) have correlation values of 0.0161, and the largest dots (Panel D) have correlation values of 0.3301. Larger dots suffer from less leakage and are thus more desirable.

F). From these results, we can determine which spherical harmonics to weigh our analysis towards—these will be the spherical harmonics which boast the largest dots overall in both leakage matrices.

In particular, we note that the size of the dots is often at a maximum for modes where $\ell = \pm m$, and therefore, it may be beneficial to prioritize our analysis to these modes, while also being cautious of potential results or detections which occur for modes with smaller dots in the leakage matrices, and thus, which suffer more from spatial aliasing. Additionally, in a practical sense, it is beneficial to prioritize analysis towards hemispherically *antisymmetric* modes (as shown circled in Figure 3.9), as these hold the capability to cancel out global noise signals (guiding, image registration, etc), and thus, should significantly decrease our noise floor in analysis of these antisymmetric modes (Gaulme et al. 2011).

Theoretically, we have the capability to examine any and all spherical harmonics of interest for the PMODE dataset, up to and including the $\ell = 100$, as this matches the spatial resolution (averaging 200 pixels across the Jovian disk) for our experiment while maintaining Nyquist sampling. However, this is assuming no degradation due to seeing-related effects. Our spatial resolution for PMODE is 0.25 arcseconds per pixel, and our average seeing for the observational campaign was 0.85 arcseconds (calculated from Table 3.1).

During our observational campaign, the apparent angular diameter of Jupiter ranged from 46.79 to 44.44 arcseconds, for an average of 45.62 arcseconds. Combining these two values, we can calculate our effective seeing-limited number of unique resolution units across

the Jovian disk:

$$\frac{45.62''}{0.85''} = 53.66 \quad (3.14)$$

Therefore, we have 53 unique and independent elements of resolution across the Jovian disk when accounting for blurring induced from the average seeing. To satisfy Nyquist sampling, we divide this value by 2 to calculate which spherical harmonic matches the practical seeing-limited spatial resolution of our instrument, which produces a (rounded down) value of $\ell=26$, as opposed to the theoretical matching spherical harmonic of $\ell=100$. Importantly, on nights and frames where the seeing value is higher than this, the maximum possible value of $\ell=26$ will decrease—thus, it may be beneficial to prioritize analysis towards the lower-order modes which will be less affected by seeing-induced variations.

However, even a value of $\ell=26$ will produce 729 separate power spectra to analyze, as we must include every m value, ranging from $-\ell$ to $+\ell$. For the sake of choosing a practical set of data to analyze over the course of this dissertation, we choose to restrict the set of modes even further—a decision which we believe to be appropriate, as again, models predict that any higher order modes will likely have lower amplitudes when integrated over a significant portion of the Jovian disk (Wu & Lithwick 2019).

Although determining which modes have the largest power distributed within them depends significantly on the spherical harmonics (convection and impacts yielding the largest power distributed in lower order modes, storms yielding the largest power distributed in higher order modes), overall, it is most often expected that a higher power will be distributed in the lower-order modes (Markham & Stevenson 2018). Finally, the previous

tentative detection of the global modes of oscillation was discovered in the $Y_\ell^m = (1, 0)$ and $Y_\ell^m = (1, \pm 1)$ modes (Gaulme et al. 2011). Combined, all these factors biases our analysis towards the lower-order modes.

The next step which remains in our spherical harmonics is to determine the highest mode that will be considered in our analysis—as we expect the highest power to be in the lowest order modes and decrease in amplitude as we increase in mode number, the next step is to determine how far from these lower order modes we should analyze as well. For this, we empirically choose the $Y_\ell^m = (11, \pm 11)$ mode—although theoretically, power distributed in this mode should be low, there is also a chance that our particular instrument will hold a higher sensitivity towards this mode.

The prime justification for interest in this particular mode is the width of the nodes of oscillation—one node of this spherical harmonic matches *half* the width of our absorption line most closely. This sizing comparison is shown in Figure 3.14. As a result of this, the entire width of our absorption line contains two nodes of the $Y_\ell^m = (11, \pm 11)$ harmonic, and thus, holds the potential to measure the *full* amplitude of any oscillation present here, from the full minima of an oscillation to its full maxima, as both protruding and receding nodes of the oscillation are entirely encompassed. Therefore, it seems a suitable choice to extend our analysis to this mode of interest as a natural cutoff.

To conclude this analysis and gain some knowledge as to which modes we should prioritize our analysis towards, we generate the spherical harmonic leakage matrix for all modes up to and including the $Y_\ell^m = (11, \pm 11)$ mode, which is shown in Figure 3.15. As in the matrices




Figure 3.14 Left panel: Average Dopplergram for the night of 06 August 2020, with basic data reduction processes (bias-, dark-, flat-, and leak-correction) applied, and a green rectangle plotted *via* SAOImage DS9 (Joye & Mandel 2003). Right panel: the $Y_{\ell}^m = (11, -11)$ spherical harmonic, calculated and normalized as described in Sections 3.2.2 and 3.4, and scaled to the same width as the left-panel Dopplergram. The same green box has been plotted over the same coordinates as in the leftmost panel, displaying that it encompasses both the dark and light nodes of the oscillation within the inscribed area.

shown in Figure 3.13, this matrix also suggests that the $\ell = \pm m$ modes suffer less spatial leakage, and are therefore preferred for analysis.

We know that the amplitudes of the modes themselves should decrease as we increase in ℓ (this is not shown in the figure, as it displays solely the spatial leakage suffered by our instrument), so we should also bias our analysis towards the lower order modes. Finally, we know that the hemispherically antisymmetric modes (for example, the $Y_{\ell}^m = (1, 0)$, $Y_{\ell}^m = (2, -1)$, $Y_{\ell}^m = (2, \pm 2)$, and so on, as shown in Figure 3.9) will display a lower noise floor due to the antisymmetric cancellation of global signals, so we will also focus our analysis towards these in particular.




Figure 3.15 Spherical harmonic leakage matrix for the PMODE observational campaign, ranging from the $Y_\ell^m = (0, 0)$ to the $Y_\ell^m = (11, \pm 11)$ spherical harmonics. Calculated for the central night of the observational campaign, 18 August 2020. The leakage matrix is obtained by calculating the spherical inner product of each spherical harmonic with itself (following isolation to the respective region of the sphere). The size of the dot corresponds to the percentage of the recovered spherical harmonic, where the smallest dots have correlation values of 0.0366, and the largest dots have correlation values of 0.3216. Larger dots suffer from less leakage and are thus more desirable. No atmospheric degradation has been included in this calculation.

In summary, we believe it is most efficient to analyze all modes up to and including $Y_\ell^m = (11, \pm 11)$, while focusing the majority of our time and efforts on modes that are: low in order, having values of $\ell = \pm m$, and hemispherically antisymmetric. Although no modes fit all of these criteria (in particular, we cannot have a mode that has both $\ell = \pm m$ and is hemispherically antisymmetric), we can bias our analysis to modes which fit any combination

of these factors. This produces an optimal subset of modes, which includes:

- All modes up to $\ell = 3$ and all corresponding m values. **Here, we have 16 amplitude spectra satisfying this requirement.**
- Each mode with $\ell = \pm m$ beyond the lowest order modes that already being considered, to include $Y_\ell^m = (4, \pm 4)$, $(5, \pm 5)$, $(6, \pm 6)$, $(7, \pm 7)$, $(8, \pm 8)$, $(9, \pm 9)$, $(10, \pm 10)$, and $(11, \pm 11)$. **Here, we have 16 additional $\ell = \pm m$ amplitude spectra satisfying this requirement.**
- Hemispherically antisymmetric modes. However, even this is a large subset of modes: the $\ell = 0$ set has 0 hemispherically antisymmetric modes, the $\ell = 1$ set has 1 hemispherically antisymmetric mode, the $\ell = 2$ set has 2 hemispherically antisymmetric modes, and so on. Thus, we have 66 hemispherically antisymmetric modes for the full set. We will limit our included plots to show the lower order of these, up to $\ell = 5$, to include $Y_\ell^m = (4, \pm 3)$, $Y_\ell^m = (4, \pm 1)$, $Y_\ell^m = (5, 0)$, $Y_\ell^m = (5, \pm 2)$, and $Y_\ell^m = (5, \pm 4)$. These are all shown circled within Figure 3.9. (We note that the $Y_\ell^m = (1, 0)$, $(2, \pm 1)$, $(3, \pm 0)$, and $(3, \pm 2)$ modes also satisfy this criteria, however, are also already included in our set of modes to analyze because they have an ℓ value lower than 3.) **Here, we have 9 additional hemispherically antisymmetric amplitude spectra satisfying this requirement.**

This final set of modes to bias our analysis towards includes 39 individual amplitude spectra. In our continued analysis, we will report the noise values for each amplitude spectrum

up to and including $\ell = 11$ and all corresponding m values, but will include only the spectrum itself for this set of predefined modes to which we expect the highest sensitivity, and any additional amplitude spectra included in the $Y_\ell^m = (0, 0)$ to $Y_\ell^m = (11, \pm 11)$ range that may display interesting features through careful analysis. This method of reporting prevents the author from including 144 separate amplitude spectra plots within this dissertation, the addition of which would certainly make this manuscript needlessly lengthy.

3.7 Additional Noise Sources

3.7.1 *Tracking and Image Registration*

Fine guiding for the PMODE observational campaign was done via a fast steering mirror operating at 0.5 KHz. It is presumed that these corrections, when integrated over a 30s interval, do not significantly contribute noise to the signal. However, Jupiter's position is only estimated via the limb finding routine, which we know to be accurate to at least 1/20th of a pixel.

These two errors are intrinsically combined and cannot be disentangled. This frame-to-frame image registration stability, with its error of approximately 0.04 pixels, provides a rough per-image noise level of $50 \text{ km s}^{-1} \times 0.04 \text{ pix} / 200 \text{ pix} = 10 \text{ m s}^{-1}$ approximately, which contributes to the total noise. This global error (induced from both tracking and registration) affects all pixels and is assumed to have no periodicity.

This registration also encompasses tracking noise contained within it, as the tracking errors will manifest as changes in the location on the detector, which is compensated for

within the image registration process. However, as we utilized tip/tilt correction on a telescope with high sensitivity and fast tracking rate, we can safely assume that the tracking noise is a small portion of the registration noise it is contained within. When including all 15,354 images from the observational campaign, this noise level reduces to $\sim 8 \text{ cm s}^{-1}$ for the final amplitude spectra.

3.7.2 Distortion Noise

Unfortunately, no distortion calibration frames were collected during the observational campaign, rendering us unable to fully calibrate any distortion-induced noise. Our requirement during instrumentation build and calibration was to produce a distortion smaller than 1/10th of a pixel, and it is believed that this goal was achieved during alignment. Nonetheless, a distortion of this magnitude could still produce spurious effects on the velocity, which may be seen in the final images shown within Section 4.2.

3.7.3 Temperature Noise

Temperature Noise

The temperature of the cell was controlled to $\pm 0.1 \text{ C}$. The experiment did not include logging of the temperature variation over time throughout the observational campaign, though quality-check routines were in place to alert the observers should temperatures exceed this acceptable variance of $\pm 0.1 \text{ C}$.

However, Cacciani et al. (2001) and Tomczyk et al. (1995) show that the error induced by temperature fluctuations on this level would be on order 10^{-3} to 10^{-4} m s^{-1} in the frequency

regime considered here. Therefore, we assume that noise induced by thermal fluctuations is significantly lower than other noise sources, and is not a point of particular concern.

CHAPTER 4

ON SIGNIFICANT RESULTS FROM THE PMODE CAMPAIGN

The rich dataset collected through PMODE on the AEOS 3.6m telescope has provided a wealth of information. From this dataset, we have validated our instrument’s velocimetric sensitivity through a sensitive measurement of the Earth’s rotational signal within our images, produced the second Doppler measurement to ever be collected of Jupiter’s zonal wind profile, and conducted thorough time-domain analysis of all modes up to and including $\ell=11$, with all corresponding m -values, detecting no significant excess signal within any of these amplitude spectra. These results hold exciting implications and constraints for the fields of Jovian atmospheric dynamics and seismology

4.1 Detection of Earth’s Rotational Signal within the PMODE Data

Prior to investigating any significant results from the PMODE campaign, we first wish to validate that our Doppler velocimeter does, in fact, have the capability to measure true velocimetric signals on-sky. We explore this validation through a search for the Earth’s rotational signal within our collected data, a signal which should not only be present, but should have an amplitude of up to 460 m s^{-1} . Measurement of this signal is still a difficult one, as it was not achieved by the most recent attempts to search for Jovian oscillations (Gaulme et al. 2011), and if detected, will provide strong confidence in both our instrument’s capabilities, and in our data reduction procedures.

As described in Section 3.2.1, this signal should appear as a predictable and gradual variation in intensity throughout the night. The Earth’s rotation over a night is known well,

and can thus be plotted over the exact nights of our observational campaign. Here, it will appear as a sine curve, gradually increasing and decreasing in velocity as we, the observer, rotate towards and away from Jupiter.

As this change appears as a global signal, no spherical harmonic multiplication is required to isolate it. As such, we conduct all data reduction techniques in Chapter 3 prior to this step. We carefully calibrate each image with our MOF unit conversion, and then simply take the average of all pixels in each image to obtain our overall, average, integrated intensity for that image. This produces one single integrated intensity time series for the duration of the observational campaign, with breaks in observation due to daytime periods preserved. Finally, we repeat this exact same process, but now in a velocity insensitive region of the Jovian disk, on the right side and far from the absorption line, to produce a second set of “velocity insensitive continuum residuals” for comparison with our velocimetric signal.

We then simply over-plot these sets of integrated intensity residuals on the sinusoidal curve of the Earth’s rotation. Both time series, for the velocity sensitive region, and a comparison with the insensitive continuum portion of the disk, are displayed in Figure 4.1. For ease of viewing, this plot has been isolated to a smaller subset of days in the observational campaign, and each time series displays a zoom-in for one representative night inset on both. For completeness, the entire velocimetric full time series is shown in Figure 4.2, although structure is harder to discern within this plot.

From Figures 4.1 and 4.2, the Earth’s rotational signal is clearly visible in the velocimetric residuals, with overall structure matching the sinusoidal structure of the Earth’s rotational

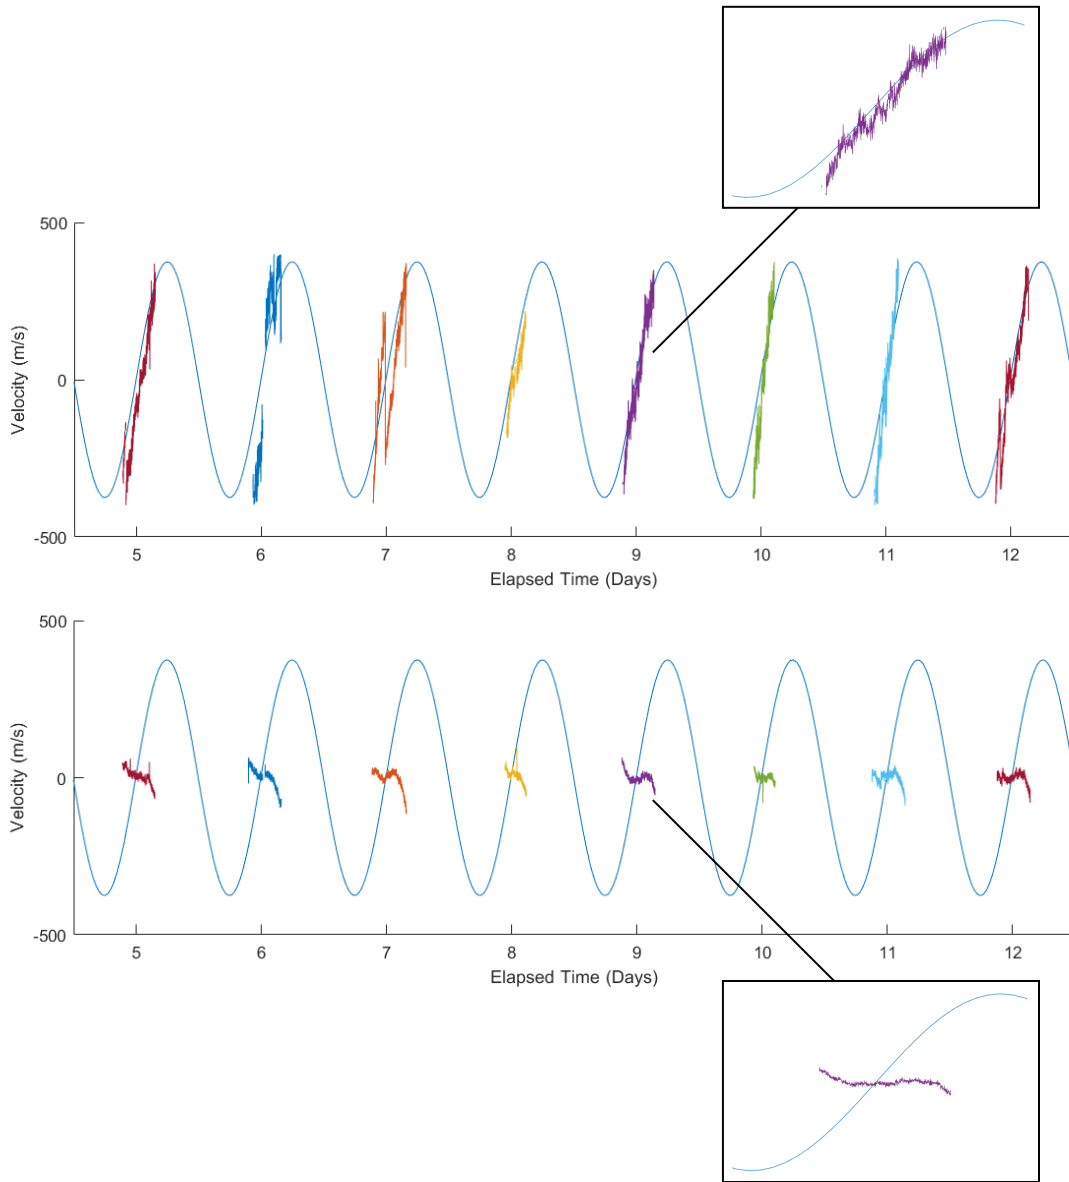


Figure 4.1 Measurement of the Earth's rotational signal in velocimetric data from the PMODE campaign. The top panel represents average integrated intensity from the velocity-sensitive absorption line regime. The bottom panel displays the corresponding average integrated intensity from an insensitive regime on the far side of the Jovian disk from the absorption line. In both, the Earth's rotational curve is plotted as a blue line, and the residual integrated intensity is over-plotted above it. These plots display a selected subset of nights from the middle of the observational campaign for ease of viewing, and both panels display a zoomed-in plot of a representative night.




Figure 4.2 Measurement of the Earth’s rotational signal in velocimetric data from the PMODE campaign, as in Figure 4.1. This plot displays the average integrated intensity for the velocity sensitive portion of the Jovian disk for the entire observational campaign as opposed to a selected central subset, for completeness. Again, the Earth’s rotational curve is plotted as a blue line, and the residual integrated intensity is overplotted above it.

curve nearly exactly, in both structure and amplitude, over the entire observational campaign. In contrast, the insensitive “continuum residuals” display no such structure or amplitude (although they maintain a significantly higher S/N, as the counts in this regime are not within the absorption line, and are thus much higher). Therefore, we proceed with confidence that our set of novel instrumentation is indeed capable of collecting on-sky Doppler velocimetric measurements.

4.2 A New Measurement of the Jovian Zonal Winds

Our first prime result from the PMODE campaign manifests as an independent, spatially resolved Doppler measurement of the Jovian zonal wind profile, only the second to ever be collected (Gonçalves et al. 2019).

This measurement, like our ultimate goal of detection of the global modes of Jupiter, is

a uniquely difficult one to conduct. Specifically, the winds themselves sustain a maximum value of approximately $\sim 0.1\%$ of that of the overall velocimetric signal due to the Jovian rotation. This requires a particularly sensitive set of instrumentation, and exceptionally careful analysis to disentangle true wind signals from all complicating, superimposed signals.

4.2.1 Extracting the Zonal Wind Profile from PMODE Data

For our purposes, we search within our average Dopplergrams—those which are implemented in the data reduction processes to remove all external velocimetric effects aside from those due from oscillatory modes. Included in these are any velocimetric signal due to atmospheric dynamics and winds. To disentangle measurement of the zonal winds from all complicating signals, we choose to apply a line-by-line smoothing filter, which we then subtract from every individual image, leaving only the components with higher spatial frequencies.

The unfortunate effect of implementing this technique to isolate signals due to the higher-frequency zonal winds is that all lower-spatial-frequency components (including the Jovian rotation and the overall absorption line profile, but also any lower-frequency-components of the ZWP with a width larger than that of our filter), are entirely removed from the final data product.

The technique used is outlined in Reach et al. (1997); we summarize the process as applied to our data here. We begin by transforming the nightly averaged, PMODE derived velocigram into a planetographic coordinate system. Next, a Fourier transform is applied to the planetographic image to create $\mathcal{F}(I)$. We then create two maps that are representative of separate frequency regimes (Reach et al. 1997): Choice of this technique provides an

avenue to not only separate undesired signals, but also as a way to differentiate zonal flows by wave-number.

For the low-frequency signals:

$$I_{\text{lf}} = \mathcal{F}^{-1} \left[\mathcal{F}(I) e^{-f^2/2f_s^2} \right] \quad (4.1)$$

To suppress high-frequency components in the original image we smooth *via*:

$$I' = \mathcal{F}^{-1} \left[\mathcal{F}(I) \left(1 - e^{-f^4/f_s^4} \right) \right] \quad (4.2)$$

This allows us to separate frequency regimes by $f_{\text{lf}}^{-1} < \theta < f_s^{-1}$ as f is the spatial frequency projected on the disk. The final computation can be described using $I_{\text{filtered}} = I' - I_{\text{lf}}$.

Planetary zonal wind profiles can be described by two primary, different regimes, because giant planets have a particularly long timescale for long-scale friction, τ_{fr} (Golitsyn (1973), Flasar (1989)), which can be understood as the frictional relaxation time, or the required time for an atmospheric perturbation to return to equilibrium. This provides a way to measure the inertial response of an atmosphere following exposure to external effects (Flasar 1989).

The respective Gas Giant zonal wind profiles can thus be decomposed into one regime which is dominated by wave-numbers smaller than that of the large-scale friction (or n_{fr}), and a second portion which contains all remaining wavenumbers greater than n_{fr} . For the case of Jupiter, in particular, n_{fr} is reported to be ~ 20 (Galperin et al. 2001). Modes below n_{fr} will be overwhelmed by the Jovian equatorial jet, and the higher modes will contain information

pertaining to the more finely detailed structure. As this smaller scale structure contains the most interesting information on the Jovian zonal winds—the interactions between the Jovian belts and zones—and it is necessary to remove all structure with a width equivalent to the absorption profile and larger, we will proceed forward by choosing this value of $n_{fr} = 20$ for our filtering.

Figure 4.3 displays a visualization of this simple applied filtering procedure, showing a Dopplergram before and after this filtering, and corresponding projections of the profile resulting from a vertical cut across the disk for both cases, prior to and following the filtering. Although it is desirable to measure the low-frequency components as well, the complexity of disentangling the absorption line signal and velocimetric effects to do so are more prevalent than anticipated, and thus, this remains for future work.

Line-of-sight velocity correction has been applied to each of these Dopplergrams, in addition to the overall MOF Unit velocimetric calibration. In the lower-left panel of Figure 4.3, which displays the high-pass filtered Dopplergram, we now see alternating dark and light bands in the region of the absorption line profile. Importantly, we see that this signal *changes in sign* when crossing the deepest portion of the absorption line—an effect that will only occur for true velocimetric signals, due to the change in gradient direction of the absorption line profile when crossing the bottom of the line, as detailed in Section 3.3.

We calculate this profile for each row in the absorption line, average each of those to produce a nightly profile, and finally, average each overall profile for every night to produce a final, overall Doppler measurement of the ZWP of Jupiter for the duration of the

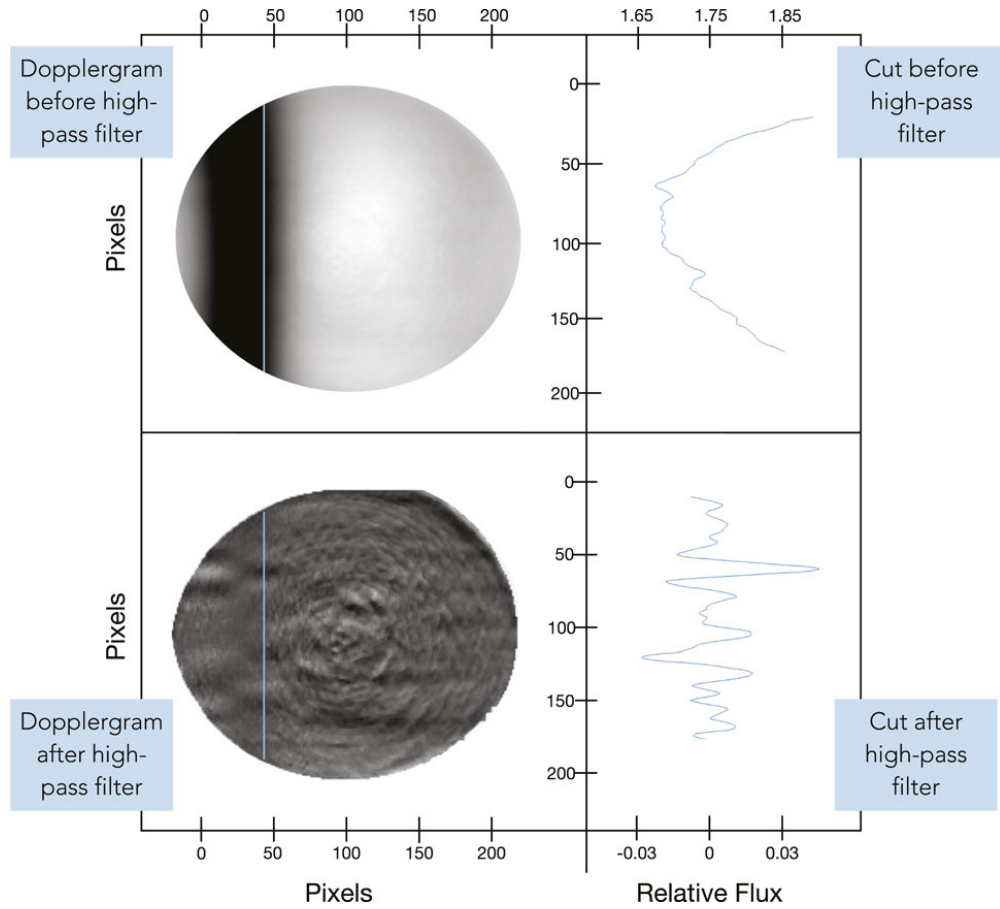


Figure 4.3 Top left panel: An averaged Dopplergram for a single night from the PMODE campaign, prior to application of high-pass filtering. Top right panel: The projected vertical cut across the Dopplergram disk, from the area marked by the blue line in the top left panel; the absorption line profile is clearly seen. Bottom left panel: The same as the above panel, after the described high-pass filtering has been applied to remove overall disk structure and the absorption line profile; the zonal winds appear clearly now as alternating dark-and-light bands within the sensitive region of the disk. Bottom right panel: The projected vertical cut across the filtered Dopplergram disk, from the area marked by the blue line in the bottom left panel; the absorption line profile has been removed and the zonal wind structure becomes apparent. Each plot has been normalized for ease of viewing, corresponding with the axis labels.

observational campaign.

4.2.2 Final PMODE Measurement of the Zonal Wind Profile

The final ZWP from our dataset is shown in Figure 4.4, with corresponding error bars calculated as the standard deviation of all averaged profiles which are combined to produce this result.




Figure 4.4 Median, high-pass-filtered and calibrated, velocimetric Jovian ZWP from the PMODE dataset. The final profile is displayed as black line, with corresponding error bars overlaid in transparent grey. Within this figure, the prograde direction of flow is represented as a positive velocity, and the retrograde direction of flow is represented as a negative velocity. A few locations of significance in this plot are marked as vertical, labeled black lines. These correspond to the Great Red Spot (GRS), the inter-sectioning region of the Northern Equatorial Belt (NEB) and Equatorial Zone (EZ), and the inter-sectioning region of the Northern Temperate Belt (NTB) and the Northern Tropical Zone (NTrZ).

The locations in Figure 4.4 that display the highest nightly variance are marked with vertical black lines, and are as follows:

First, our profile displays increased variance within the region of the Great Red Spot (GRS), at -20° S. This is expected, as the GRS rotates in and out of our sensitivity region throughout the course of the night, and should contribute an increase in radial upwelling to the zonal wind signal when it is included, and a decrease when it is excluded, resulting in an overall increased change around the median value.

Secondly, we note increased variance in our ZWP at the 7° N region. This latitude falls at the transitory region between the NEB and Equatorial Zone (EZ), a region of the disk which is expected to host a trapped, planetary-scale Rossby wave. This wave is theorized to be the source of known upwelling and hot-spots located at this latitude (Barrado-Izagirre et al. (2013), Arregi et al. (2006)), features which would certainly increase the variance in this regime if present, and which may be visible in time-series analysis as well as an increase in low-frequency power.

Finally, we have highlighted the intersection of the NTB and Northern Tropical Zone (NTrZ) at the $\sim 21^\circ$ N region as an area of significantly increased variance in our measurement. This region is known to display year-to-year variation, with significantly higher outbreaks occurring every five years (Sánchez-Lavega et al. 2017). As with the NEB/EZ region, these features are also expected to increase variance within our data, which our signal agrees with.

One important note is that the variations in structure which we are attributing to ve-

locimetric measurements of the ZWP do not correlate with the locations of belts and zones present on the planet. Thus, we can have confidence that this alternating structure is not simply due to changes in intensity on the Jovian disk due to its intrinsic structure, but rather to true velocimetric measurements. This is shown in Figure 4.5, which plots our measured ZWP upon a representative continuum image from the PMODE observational campaign.




Figure 4.5 The median, high-pass-filtered and calibrated, velocimetric Jovian ZWP from the PMODE dataset, as shown in Figure 4.4. Here, the profile is overlaid on a representative continuum image collected on 15 August 2020 from the PMODE campaign. We note that the changes in the zonal wind velocities are not correlated with the alternating belt-zone structure, and thus can be assumed to be attributed to true velocities as opposed to simple changes in intensity.

4.2.3 Comparison of PMODE ZWP with Historical Measurements

As a reminder from Section 1.5.1, we note that only one other velocimetric measurement of the Jovian ZWP has ever been collected, by Gonçalves et al. (2019). Importantly, this velocimetric measurement displayed significant differences from the cloud-tracking profile, which is extremely well-characterized, but maps only the overall large scale velocity as opposed to the true particle velocity.

The differences from the Gonçalves et al. (2019) and the cloud-tracking profile, primarily in the 0–15° N region, imply that the true particle velocity of the zonal winds significantly differs from the large-scale motion of the winds. This discrepancy could have significant changes on Jovian interior models and constraints on the interior structure, as the most recent predictions (Militzer et al. 2022) allow the ZWP to vary to best fit the measured gravitational moments of Jupiter; this model implies that Jupiter could produce the measured gravity if entirely lacking a solid core. Placing a defined constraint on the true velocities of these winds will provide an avenue for further refined analysis on the current Jovian atmospheric and interior models.

Because of the difference between the sole previous velocimetric measurement and the feature-tracking measurement, we are particularly interested in comparing our new measurement with the previous Doppler measurement. These profiles are shown together in Figure 4.6; the Gonçalves et al. (2019) profile has been simply approximated from their 2015 measurements. The true profile as displayed in their original paper includes estimated error and a second profile from their 2016 measurements, neither of which is displayed in Figure

4.6. We refer the reader to the full Gonçalves et al. (2019) paper for a more accurate and thorough representation of the previous measurement of the ZWP.




Figure 4.6 The new PMODE Doppler measurement of the Jovian ZWP, shown in black, compared to sole previous Doppler measurement of the ZWP (approximated and adapted from Gonçalves et al. (2019), with permission, the true profile may vary slightly), shown in blue. There is a difference in resolution, and in structure, particularly within the 0–15° N region. These profiles were measured at different years and different atmospheric scale heights.

From Figure 4.6, we see that these two separate velocimetric measurements of the Jovian ZWP display some substantial differences. First, there is certainly a difference in spatial resolution, of which the PMODE data has an increase by a factor of nearly 3; however, we look beyond that at the true overall structure and amplitudes of the two profiles. Our noise levels

at higher latitudes rise, due to complications with line-of-sight velocity falloff corrections and boundary effects from the high-pass filtering. For the purposes of this comparison, we will consider only the mid-latitude regimes, from -30° to $+30^\circ$. Within this region, the two profiles are structurally similar from -30° to -10° , at which point they begin to diverge. The previous ZWP measurement (Gonçalves et al. 2019) reaches higher overall velocities than the PMODE measurement, displaying a maximum of $\sim 80 \text{ m s}^{-1}$ at in the -20° to -10° region, as opposed to our profile, which reaches a maximum of 33.15 m s^{-1} in the same region.

However, the prime discrepancy can be found in the 0° to $+15^\circ$ region. Here, the Gonçalves et al. (2019) profile displays a clear local minima, while the PMODE profile reaches a local maxima. We note that that Gonçalves et al. (2019) state that in this region, their profile significantly differs from the well-studied cloud tracking profile in the same way, where the Doppler measurement reaches a minimum where the cloud tracking reaches a maximum, implying that the true particle velocities are much lower than that suggested by cloud tracking. Beyond this regime, both Doppler measurements again reach agreement in structure at the NTB/NTrZ interface, where they both reach a local maximum.

These variations are interesting, and subsequently beg the questions: what are the true velocities present in the 0° to 15° N latitude regime? Secondly, why would two independent Doppler measurements display opposition to each other within this band? It is certainly true that the two Doppler measurements were collected nearly six years apart, and physical properties of the winds may have changed on the planet; however, we know the winds to historically be exceptionally stable over timescales of this length (García-Melendo & Sánchez-

Lavega 2001), so we believe this to be unlikely. Secondly, these measurements were collected in two different wavelengths (517 nm (Gaulme et al. 2008) vs 770 nm), and thus two different atmospheric heights, and the ZWP may vary with depth. Although this is a feasible explanation for some discrepancies, we believe that the difference between these two measurements is likely due to some other source.

On our journey to search for the cause of this discrepancy, we will next compare our PMODE Doppler measurement of the ZWP with the well-studied cloud-tracking ZWP. The *Hubble* measurements for comparison were obtained from the Outer Planet Atmospheres Legacy (OPAL) database (Simon et al. 2015). This comparison is shown in Figure 4.7. Here, we see immediately that the structure of our Doppler measurement with the cloud-tracking measurement matches exquisitely well over the full mid-latitude regime, and similarities are even shown to higher latitudes, with well-matching structure ranging all the way to -60° in the Southern hemisphere. This overall structural similarity with the cloud-tracking profile from the Hubble Space Telescope (HST) gives us a great deal of confidence in the validity and accuracy of our Doppler measurement.

However, there are two glaringly obvious differences between the PMODE Doppler ZWP and the *Hubble* Cloud Tracking ZWP as shown in Figure 4.7. First, a significant discrepancy is present in the equatorial regime of the ZWP—the HST profile displays a strong prograde velocity of 78 m s^{-1} , while our PMODE measurement displays a retrograde velocity of 18 m s^{-1} , implying that the winds are flowing in two different directions in the respective measurements.

Secondly, a drastic difference is seen in the amplitudes themselves, which differ by as much as 115 m s^{-1} in the -10° regime. The large discrepancies in amplitude, and strong similarity in overall structure, is particularly curious. Now, we return to an important consideration in our measurement of the ZWP—to obtain it, we have to filter out *all low frequency components*. It is certainly possible that *all* the missing amplitude in our measurement could be present in these low frequency components themselves.




Figure 4.7 The new PMODE Doppler measurement of the Jovian ZWP, shown in black, compared to the well-studied cloud-tracking ZWP measured by the HST, shown in blue. The HST data was obtained from Simon et al. (2015). It is clear that there are differences in amplitude, but the overall structure is similar.

To determine the validity of this hypothesis, we turn to past work. Again, we return to

Galperin et al. (2001), who display clearly how the cloud-tracking ZWP can be decomposed into low- and high-frequency components through Legendre-polynomial decomposition; this can be thought of as a tool to simply separate low- and high-spatial effects on the zonal winds. To validate this, we take the HST ZWP as provided by the OPAL database, and decompose it into Legendre polynomials as reported by Galperin et al. (2001). This is visualized in Figure 4.8A, where we decompose the full cloud-tracking profile (plotted in blue) into the low-order Legendre polynomials below $P_\ell^m = 20$ (plotted in green), and the remaining higher order-polynomials necessary to produce the original profile (plotted in orange.) Both the decomposed profiles themselves are shown in Figure 4.8B; these, when directly summed, reproduce the exact HST cloud-tracking profile.

Interestingly, the higher-order Legendre polynomials (orange in Figure 4.8B) display a striking similarity to the PMODE ZWP measurement—in particular, the high order components of the cloud tracking profile *also* display a retrograde velocity at the equator, and comparable amplitudes to our measurement. To directly compare between our filtered Doppler measurement and these components, we simply over-plot our measured ZWP over this Legendre decomposition, which is shown in Figure 4.8C. We note that the high-order Legendre decomposition of the HST ZWP (plotted in orange) and our PMODE ZWP (plotted in black) match particularly well. There is some appearance of phase shift difference between the orange curve—we acknowledge that we are averaging multiple nights and multiple longitudes to produce our final zonal wind profile, which may induce some overall variations in the final reference frame. Even more interestingly from this first similarity, we note also that

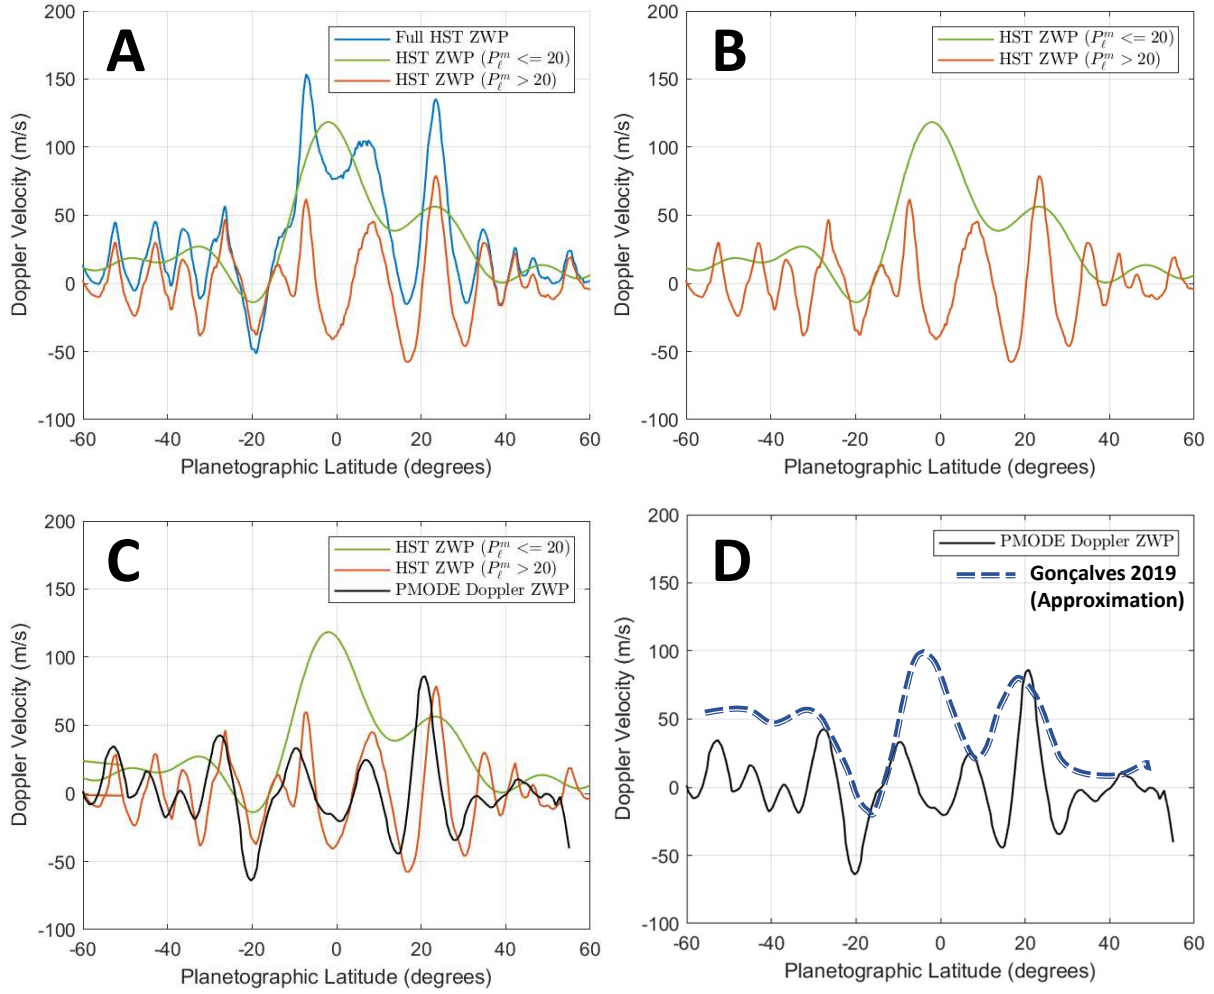


Figure 4.8 Panel A: The HST ZWP from the OPAL database (blue), decomposed into Legendre polynomials $P_\ell^m \leq 20$ (green) and $P_\ell^m > 20$ (orange). Panel B: The low- and high-order Legendre decomposition from the HST ZWP (green and orange respectively), shown individually for comparison with Panel D. Panel C: The low- and high-order Legendre decomposition from the HST ZWP (green and orange respectively), with the PMODE Doppler ZWP measurement overplotted in black. We note a strong similarity between the orange line and the black line in this panel. Panel D: The PMODE ZWP measurement (black) with the approximated Gonçalves et al. (2019) Doppler ZWP measurement (dashed blue). We note the similarities between this panel, and Panel B.

the previous Doppler measurement of the ZWP (Gonçalves et al. 2019) has some significant similarities to the low-order Legendre decomposition of the HST ZWP. We have overplotted

our approximation of their measurement over the PMODE measurement in Figure 4.8D, for comparison with Figure 4.8C.

4.2.4 New Implications for Jovian Wind Velocities from the PMODE Measurement

The strong similarities between Figures 4.8C and 4.8D holds an exciting implication for Jovian atmospheric science. Namely, we believe that the Gonçalves et al. (2019) measurement did not measure the entire ZWP, and was instead sensitive only to the low-order regime dominated by the equatorial jet. Therefore, we believe the discrepancy between the Gonçalves et al. (2019) Doppler profile and the cloud-tracking ZWP was not a true discrepancy—rather, they were simply missing a significant piece of the puzzle.

Likewise, this implies that the PMODE Doppler profile is not sensing the entire ZWP either, and that the low-order regime which is excluded by the high-pass filtering contains our missing piece of the puzzle, as well. This suggests that, if the previous Doppler measurement is sensitive only to the low-order spatial frequencies, and the PMODE Doppler measurement is sensitive only to the high-order spatial frequencies, a combination of the two Doppler profiles may *exactly reproduce* the well-studied cloud tracking ZWP. Therefore, the intrinsic velocity of the small-particles, and the overall velocity and the large-scale features may in fact be similar, instead of differing substantially as believed from the Gonçalves et al. (2019) measurement. This is significant, as an accurate measurement of the true particle velocities is one of the largest constraints for creating an accurate Jovian atmospheric and interior model for further determination of the true interior structure.

Unfortunately, we do not have access to the Gonçalves et al. (2019) dataset to directly test this theory. However, as we believe this profile is well-represented by the low-order Legendre polynomial fit to the HST ZWP (as from the comparison in Figures 4.8C and 4.8D), we will instead use the low-frequency components of the cloud tracking profile to combine with our measurement, effectively replacing the large spatial frequencies that were filtered out in our data analysis process. The goal of this test is to determine if a combination of a Doppler high-spatial frequency measurement with a separate low-frequency measurement will accurately reproduce the full and well-studied cloud tracking profile. To determine this, we simply add these two curves—those represented by green and black in Figure 4.8C. The result of this addition is shown in Figure 4.9.

Immediately from Figure 4.9A, we can begin to see the well-known structure of the large-scale feature-tracking ZWP. For ease of comparison, we plot the combination of the high-order PMODE data with the low-order HST data over the *full* HST feature-tracking ZWP. Both profiles overplotted can be seen in Figure 4.9B. Following combination of our new and sensitive Doppler measurement with the low-order components of the cloud tracking profile, we reproduce a very similar profile to the full HST ZWP, not only in structure, but also in amplitude. There are some slight differences, which we primarily attribute to our longitudinally-averaged zonal wind profile throughout multiple nights, likely inducing variance in our final reference frame.

However, the strong overall similarity shown in this comparison gives us confidence in our final claim that *neither* our PMODE profile nor the sole previous Doppler measurement




Figure 4.9 Panel A: The PMODE ZWP, which we know to contain only the high-spatial frequency components, added to the low-spatial frequency components of the HST cloud-tracking ZWP. Panel B: as in Panel A, but overplotted with the full, well studied cloud tracking ZWP from HST (in blue) for comparison.

produce a full and complete picture of the zonal winds, and each measurement is sensing a different regime of the zonal winds, which explains the differences in our two profiles. The true particle velocity of the Jovian zonal winds is a combination of both these low- and high- order sensitive measurements. We believe this discovery implies that the true particle velocity within the Jovian zonal winds accurately matches the large scale feature-tracking velocity, as opposed to the fundamental difference which was previously assumed.

4.3 Characterization of Power Spectrum Noise Levels

Before we can begin to thoroughly analyze all final power spectra, we must first ensure that we have a careful and accurate way to measure the resultant noise levels in each. For the sake of description in this section, we will use the $Y_l^m = (1,0)$ spectrum as an example, although the same techniques are applied to each mode of interest, up to and including the $\ell = 11$ modes. To characterize these noise levels of the background in our power spectra, we implemented multiple techniques, as described below.

4.3.1 “Random-Permutation” (RP) Noise Characterization

First, to attempt to accurately measure the average, overall noise level of our full, 24-night power spectra for each mode, we simply generated a permuted realization of our time series, with all daytime-gaps preserved so that any effects from the window function are included in the resultant spectrum. To do this, we read in each true nightly set of data, then smooth our data with a robust quadratic regression and a smoothing window of 200 samples, or $166.67 \mu\text{Hz}$, to remove any long-term nightly trends. We then shuffle the night of

data using the in-built **MATLAB** function `randperm`, prior to multiplication by an equal-length Tukey window with a cosine fraction of 0.1 to minimize edge effects in the Fourier domain.

Next, we apply an outlier rejection with a threshold factor of four median absolute deviations (or MAD). Each of these pre-Fourier transform steps (with the exception of the random permutation) is the exact same as for true power spectrum calculation (detailed in Section 3.5), for consistency. Finally, we then compute the Lomb-Scargle periodogram of this randomly permuted time series to produce a final “noise spectrum” with frequency resolution exactly equal to that of our true dataset. An example of this “noise spectrum” is shown in Figure 4.10. Importantly, due to the random shuffling, the low-frequency regime (up to $\sim 700\mu\text{Hz}$), does not display the higher values that we will later see in all true spectra with un-shuffled values, and thus, provides a clean average baseline noise value.

We then obtain a median value of this noise level of to estimate the average background level of noise for the entire spectrum, which is calculated on a range from 0 to $6000\ \mu\text{Hz}$. For the example $Y_l^m = (1,0)$ power spectrum, this average value is $2.03\ \text{cm s}^{-1}$, and the maximum value is $7.89\ \text{cm s}^{-1}$. All Random-Permutation (RP) background noise values for all spectra up to and including the $\ell = 11$ modes are shown in Tables 4.1 through 4.4.

4.3.2 “Background-Fitting” (BF) Noise Characterization

An additional complicating factor is that this background noise certainly changes with frequency—thus, in addition to our average value, it is also necessary to estimate the noise level specifically at the frequency range of interest. To do so, we first fit the background of each power spectrum—however, this is not trivial, because a least square fitting is not




Figure 4.10 “Random-Permutation” (RP) background noise example Lomb-Scargle periodogram for the $Y_l^m = (1, 0)$ mode, calculated for the full 24-night dataset with proper daytime observing gaps preserved.

appropriate for fitting one full power spectrum, as the distribution of a power spectrum is effectively a chi-square with two degrees of freedom.

To compensate for this, we instead fit the average of many power spectrum segments (in particular, we divide our dataset into twelve two-night segments), similarly to the process which is thoroughly described in Appourchaux et al. (2014) and Gabriel et al. (2002). Each individual time series has the same $166.67 \mu\text{Hz}$ robust quadratic regression smoothing, Tukey window multiplication, and 4-MAD outlier rejection, exactly as for the true power spectrum and the RP background noise spectrum—the only difference is that these techniques are applied a time series encompassing only two nights at a time, and this two-night power spectrum is subsequently calculated.

Once these twelve realizations have been averaged, the chi-square noise has been substantially reduced, and the background noise level can then be fit with a least square routine. In particular, we fit the noise spectrum to a two-term Gaussian (similarly to the background-fitting described in Gaulme et al. (2011)) to compensate for both the low-frequency excess noise regime, and the relatively flat regime beyond $\sim 700\mu\text{Hz}$.




Figure 4.11 Background-Fitting (BF) noise example Lomb-Scargle periodogram for the $Y_l^m = (1, 0)$ mode, calculated from 12 averaged two-night Lomb-Scargle periodograms to decrease the chi-square noise contained within the spectrum. This resultant spectrum (blue) has been fit with a least-squares two-component Gaussian (orange). Median values are obtained over differing regions of this orange curve to fully characterize the background noise.

This allows us to then calculate the median value of the background fit itself in multiple regimes—for this analysis, we divide all power spectra into three sections:

1. The low frequency regime, containing the highest noise from instrumental effects and upwelling, defined hereafter as the 0–1000 μHz regime.

2. The “regime of interest”, or where modes have previously been reported to be detected (Gaulme et al. 2011), and the area where modes are also predicted to be for nearly all theorized excitation mechanisms (Markham & Stevenson 2018), defined hereafter as the 1000–3000 μHz regime.
3. The high-frequency regime, defined hereafter as the 3000–6000 μHz regime, which maintains the lowest noise level and encompasses the highest frequencies at which models predict global modes can be found (Markham & Stevenson 2018). This region stretches to the extent of the frequencies for which we calculate our power spectra through this analysis.

All BF noise values for all spectra up to and including the $\ell = 11$ modes, calculated for the 0–1000 μHz , 1000–3000 μHz , and 3000–6000 μHz regimes are shown in Tables 4.1–4.4.

l	m	Mean Noise 0-1000 μHz (cm/s, BF)	Mean Noise 1000-3000 μHz (cm/s, BF)	Mean Noise 3000-6000 μHz (cm/s, BF)	Mean Noise 0-6000 μHz (cm/s, BF)
0	0	54	33	24	14
1	-1	42	27	20	11
<i>1</i>	<i>0</i>	6	4.5	4.2	2
1	1	39	26	19	11
2	-2	56	36	26	15
<i>2</i>	<i>-1</i>	5.7	4.5	4.1	2
2	0	36	25	18	11
<i>2</i>	<i>1</i>	5.3	4	3.7	1.9
2	2	13	9.6	7.7	4.9
3	-3	28	20	15	12
<i>3</i>	<i>-2</i>	8.8	6.4	5.9	2.9
3	-1	16	11	8.4	5.4
<i>3</i>	<i>0</i>	8.5	6.2	5.6	3
3	1	15	11	8.6	6.9
<i>3</i>	<i>2</i>	3.1	2.5	2.2	1.3
3	3	33	20	16	8.9
4	-4	22	16	13	8.4
<i>4</i>	<i>-3</i>	5.6	4.7	4.5	2.5
4	-2	13	9.9	8.5	7
<i>4</i>	<i>-1</i>	5.8	4.2	3.8	2
4	0	12	8.7	7.9	6.7
<i>4</i>	<i>1</i>	5.2	3.9	3.8	2.2
4	2	5.3	4.3	3.7	3.3
<i>4</i>	<i>3</i>	5.4	4.4	3.9	2.1
4	4	40	26	19	11
5	-5	29	19	15	8.7
<i>5</i>	<i>-4</i>	5.8	4.9	4.5	2.5
5	-3	8.5	6.9	6.4	6.1
<i>5</i>	<i>-2</i>	6.5	5.2	4.9	2.8
5	-1	7	5.5	4.8	3.1
<i>5</i>	<i>0</i>	7.1	5.4	5.2	3
5	1	8.1	6.2	5.4	5.1
<i>5</i>	<i>2</i>	2.7	2.4	2.2	1.3
5	3	6.7	5.4	4.7	2.3
<i>5</i>	<i>4</i>	7.1	5.7	5.2	2.7
5	5	22	17	14	10

Table 4.1 Average noise values for $Y_l^m = (0,0)$ to $Y_l^m = (5,5)$ power spectra from the PMODE campaign. BF signifies that the value is calculated per Background-Fitting, or a linear-least-squares Gaussian fit to 12-averaged dual-night power spectra. RP signifies that the value is calculated per Random-Permutation, or shuffling the values in each night, prior to calculation of a single 24-night “noise spectrum” with all daytime gaps preserved. Hemispherically antisymmetric modes are italicized, for reference, and all display lower noise levels than their non-hemispherically antisymmetric counterparts.

l	m	Mean Noise 0-1000 μHz (cm/s, BF)	Mean Noise 1000-3000 μHz (cm/s, BF)	Mean Noise 3000-6000 μHz (cm/s, BF)	Mean Noise 0-6000 μHz (cm/s, BF)
6	-6	20	15	12	6.4
<i>6</i>	<i>-5</i>	7	<i>5.2</i>	<i>4.5</i>	<i>2.3</i>
6	-4	10	7.5	6.7	5.7
<i>6</i>	<i>-3</i>	<i>4.1</i>	<i>3.7</i>	<i>3.7</i>	<i>2.6</i>
6	-2	8.8	6.9	6.2	4.7
<i>6</i>	<i>-1</i>	<i>5.2</i>	<i>3.8</i>	<i>3.6</i>	<i>1.9</i>
6	0	9.2	7.4	6.6	5.2
<i>6</i>	<i>1</i>	<i>4.2</i>	<i>3.6</i>	<i>3.5</i>	<i>2.2</i>
6	2	4.6	3.6	3.2	3
<i>6</i>	<i>3</i>	<i>4.7</i>	<i>3.5</i>	<i>3.3</i>	<i>1.6</i>
6	4	7.7	6.2	5.5	3.5
<i>6</i>	<i>5</i>	<i>6.3</i>	<i>5.1</i>	<i>5.2</i>	<i>2.9</i>
6	6	21	15	12	9.4
7	-7	14	12	9.2	6.1
<i>7</i>	<i>-6</i>	<i>5.6</i>	<i>4.4</i>	<i>4.3</i>	<i>2.2</i>
7	-5	11	8.6	7.3	4.7
<i>7</i>	<i>-4</i>	<i>4.9</i>	<i>4.2</i>	<i>3.9</i>	<i>2.5</i>
7	-3	7.2	5.8	5.2	4.7
<i>7</i>	<i>-2</i>	<i>5.8</i>	<i>5</i>	<i>4.7</i>	<i>2.7</i>
7	-1	4.9	4.5	3.9	2.6
<i>7</i>	<i>0</i>	<i>6.2</i>	<i>5.1</i>	<i>4.9</i>	<i>2.9</i>
7	1	6.3	5.1	4.5	4.3
<i>7</i>	<i>2</i>	<i>2.9</i>	<i>2.2</i>	<i>2.2</i>	<i>1.3</i>
7	3	6	4.4	3.8	2.1
<i>7</i>	<i>4</i>	<i>5.1</i>	<i>4.3</i>	<i>4</i>	<i>2.2</i>
7	5	7.8	6.6	5.9	4.9
<i>7</i>	<i>6</i>	<i>6.7</i>	<i>5.6</i>	<i>5.2</i>	<i>3.2</i>
7	7	17	13	11	9.7
8	-8	14	11	8.9	7.8
<i>8</i>	<i>-7</i>	<i>5.9</i>	<i>5</i>	<i>5.1</i>	<i>2.6</i>
8	-6	9.4	7.3	6.1	3.5
<i>8</i>	<i>-5</i>	<i>5.6</i>	<i>4</i>	<i>3.9</i>	<i>2.1</i>
8	-4	7.4	5.9	5.2	4.6
<i>8</i>	<i>-3</i>	<i>4.1</i>	<i>3.9</i>	<i>3.8</i>	<i>2.6</i>
8	-2	6.7	6.1	5.3	4.2
<i>8</i>	<i>-1</i>	<i>4.4</i>	<i>3.4</i>	<i>3.5</i>	<i>1.9</i>

Table 4.2 Average noise values for $Y_l^m = (6, -6)$ to $Y_l^m = (8, -1)$ power spectra from the PMODE campaign. BF signifies that the value is calculated per Background-Fitting, or a linear-least-squares Gaussian fit to 12-averaged dual-night power spectra. RP signifies that the value is calculated per Random-Permutation, or shuffling the values in each night, prior to calculation of a single 24-night “noise spectrum” with all daytime gaps preserved. Hemispherically antisymmetric modes are italicized, for reference, and all display lower noise levels than their non-hemispherically antisymmetric counterparts.

l	m	Mean Noise 0-1000 μHz (cm/s, BF)	Mean Noise 1000-3000 μHz (cm/s, BF)	Mean Noise 3000-6000 μHz (cm/s, BF)	Mean Noise 0-6000 μHz (cm/s, BF)
8	0	6.8	6.6	6.1	4.7
8	1	3.8	3.7	3.6	2.3
8	2	4.2	3.5	3.2	2.8
8	3	4.3	3	2.9	1.5
8	4	5.7	4.7	4.3	2.9
8	5	4.8	4.2	4.2	2.7
8	6	12	8.7	7.4	5.8
8	7	7.1	5.5	5	3
8	8	16	13	9.8	6.9
9	-9	18	14	11	9.5
9	-8	7.1	5.6	5.3	3.2
9	-7	8.4	7	5.9	3.8
9	-6	4.4	3.7	3.6	1.9
9	-5	7.3	5.6	5.4	4.1
9	-4	4.8	3.5	3.6	2.3
9	-3	5.9	5	4.8	4.5
9	-2	5	4.8	4.5	2.7
9	-1	4.6	4.3	4	2.5
9	0	5.5	5.2	4.9	2.9
9	1	5.3	4.7	4.6	4
9	2	3	2	2.2	1.3
9	3	4.7	4	3.5	1.8
9	4	4.3	3.9	3.8	2.1
9	5	7.1	5.8	4.9	3.8
9	6	5.8	4.7	4.6	2.9
9	7	11	8.6	7.3	6
9	8	6.7	5.2	4.7	2.5
9	9	14	11	8.7	6
10	-10	18	14	10	9
10	-9	7.2	6	5.5	3.5
10	-8	9.5	7.7	6.2	4.7
10	-7	5.7	4.5	4.2	2.4
10	-6	6.4	5.2	4.7	2.6
10	-5	4.7	3.5	3.4	1.9
10	-4	6.5	5.4	5.3	4.4
10	-3	4.1	3.8	3.6	2.6

Table 4.3 Average noise values for $Y_l^m = (8, 0)$ to $Y_l^m = (10, -3)$ power spectra from the PMODE campaign. BF signifies that the value is calculated per Background-Fitting, or a linear-least-squares Gaussian fit to 12-averaged dual-night power spectra. RP signifies that the value is calculated per Random-Permutation, or shuffling the values in each night, prior to calculation of a single 24-night “noise spectrum” with all daytime gaps preserved. Hemispherically antisymmetric modes are italicized, for reference, and all display lower noise levels than their non-hemispherically antisymmetric counterparts.

l	m	Mean Noise 0-1000 μHz (cm/s, BF)	Mean Noise 1000-3000 μHz (cm/s, BF)	Mean Noise 3000-6000 μHz (cm/s, BF)	Mean Noise 0-6000 μHz (cm/s, BF)
10	-2	6.5	5.9	5.6	4
<i>10</i>	<i>-1</i>	<i>4</i>	<i>3.5</i>	<i>3.4</i>	<i>1.8</i>
10	0	7.1	6.7	6.4	4.7
<i>10</i>	<i>1</i>	<i>4</i>	<i>3.5</i>	<i>3.3</i>	<i>2.2</i>
10	2	3.5	3.2	3.3	2.6
<i>10</i>	<i>3</i>	<i>3.6</i>	<i>2.9</i>	<i>2.9</i>	<i>1.5</i>
10	4	4.7	4.4	3.9	2.4
<i>10</i>	<i>5</i>	<i>5</i>	<i>4.3</i>	<i>4</i>	<i>2.5</i>
10	6	9.2	7.2	6.4	4.9
<i>10</i>	<i>7</i>	<i>5.2</i>	<i>4.5</i>	<i>4.4</i>	<i>2.6</i>
10	8	12	8.9	7.4	5.8
<i>10</i>	<i>9</i>	<i>5.8</i>	<i>4.9</i>	<i>4.7</i>	<i>2.4</i>
10	10	18	14	10	7.6
11	-11	24	15	12	8.7
<i>11</i>	<i>-10</i>	<i>7.3</i>	<i>5.5</i>	<i>5.2</i>	<i>3.1</i>
11	-9	12	9.6	7.5	6
<i>11</i>	<i>-8</i>	<i>5.5</i>	<i>4.8</i>	<i>4.5</i>	<i>2.9</i>
11	-7	6	5.4	4.5	2.7
<i>11</i>	<i>-6</i>	<i>3.7</i>	<i>3.5</i>	<i>3.2</i>	<i>1.7</i>
11	-5	5.8	4.8	4.8	3.4
<i>11</i>	<i>-4</i>	<i>4.3</i>	<i>3.5</i>	<i>3.7</i>	<i>2.3</i>
11	-3	6	5	4.8	4.5
<i>11</i>	<i>-2</i>	<i>5.4</i>	<i>4.4</i>	<i>4.2</i>	<i>2.6</i>
11	-1	4.6	4.1	4	2.6
<i>11</i>	<i>0</i>	<i>6.2</i>	<i>4.8</i>	<i>4.7</i>	<i>2.9</i>
11	1	5.3	4.8	4.6	4
<i>11</i>	<i>2</i>	<i>2.5</i>	<i>2.2</i>	<i>2.2</i>	<i>1.3</i>
11	3	4	3.7	3.3	1.7
<i>11</i>	<i>4</i>	<i>4.1</i>	<i>3.7</i>	<i>3.6</i>	<i>2</i>
11	5	6	4.8	4.9	3.6
<i>11</i>	<i>6</i>	<i>5.3</i>	<i>4</i>	<i>4.2</i>	<i>2.7</i>
11	7	8.5	7	6.4	5
<i>11</i>	<i>8</i>	<i>5.2</i>	<i>4.4</i>	<i>4.4</i>	<i>2.3</i>
11	9	11	8.5	7	4.4
<i>11</i>	<i>10</i>	<i>7</i>	<i>5.6</i>	<i>5.4</i>	<i>3</i>
11	11	20	13	10	8.4

Table 4.4 Average noise values for $Y_l^m = (10, -2)$ to $Y_l^m = (11, 11)$ power spectra from the PMODE campaign. BF signifies that the value is calculated per Background-Fitting, or a linear-least-squares Gaussian fit to 12-averaged dual-night power spectra. RP signifies that the value is calculated per Random-Permutation, or shuffling the values in each night, prior to calculation of a single 24-night “noise spectrum” with all daytime gaps preserved. Hemispherically antisymmetric modes are italicized, for reference, and all display lower noise levels than their non-hemispherically antisymmetric counterparts.

4.4 Final PMODE Amplitude Spectra: The (1,0) Mode

Finally, now that all careful amplitude calibration and background calibration has been completed, we look to analyze the resultant amplitude spectra for the 2020 PMODE observational campaign. First and foremost, the primary mode of interest is the $Y_\ell^m = (1, 0)$ mode, as this is the mode where previously detected global modes have been reported (Gaulme et al. 2011). The resultant power spectrum, calculated *via* the techniques described in Section 3.5, is shown in Figure 4.12.

This figure, in particular, is scaled in both X- and Y-axes to parallel the final resultant power spectrum from the Gaulme et al. (2011) paper. Importantly, we note that for this figure, we do display the *power* spectrum, which is simply the *amplitude* spectrum that we typically describe, but squared, for comparison with past results. (The classic amplitude spectrum for this mode is shown in Figure 4.13.) Here, the background is fit *via* our BF techniques as described in Section 4.3.2, and represented as a red dashed line.

We note that in this power spectrum, and all subsequent amplitude spectrum from our observational campaign that will be shown, there is an increase in power in the lowest frequency (i.e., below $700 \mu\text{Hz}$) regime. We attribute the source of this excess power to a combination of image-tracking and frame-to-frame registration errors, as discussed in Section 3.7.1. We believe this low-frequency power is in part, also attributed to true, Jovian-intrinsic upwelling events, as discussed in Section 4.2.2, that we know to be present in the 7° N regime, where the interface between the NEB and EZ is located. Regardless, we recognize that this low-order power is present (in both our dataset and all past attempts to detect the Jovian




Figure 4.12 The final power spectrum for the $Y_\ell^m = (1, 0)$ mode from the 2020 PMODE observational campaign. We note that this is the power spectrum, or our amplitude spectrum squared, for comparison with the previous reported detection of global modes within this particular Y_ℓ^m mode. The dashed red curve represents our noise floor that varies with frequency as calculated by our “background fitting” technique. We do not see excess power in the $1200 \mu\text{Hz}$ regime as previously reported with amplitudes higher than other noise peaks at higher frequencies throughout the spectrum.

oscillations), but we believe that it is attributed to observational, analysis, and intrinsic planetary noise sources, rather than attributed to any global oscillations.

The Gaulme et al. (2011) power spectrum that we are comparing to displays an excess of power reaching to a maximum of $275 \text{ cm}^2 \text{ s}^{-2}$ in the $1200 \mu\text{Hz}$ regime, which we do not see within our power spectrum. In fact, we see no significant excess power reaching any values near these, whatsoever—not in the $1200 \mu\text{Hz}$ regime, nor beyond. We note that our average noise floor in the $1200 \mu\text{Hz}$ regime of interest is $\sim 20.25 \text{ cm}^2 \text{ s}^{-2}$, while the noise floor for the previously reported detection is approaching $\sim 70 \text{ cm}^2 \text{ s}^{-2}$. With our careful data reduction and analysis techniques, we see no significant, organized power indicative of global modes within this power spectrum. We note that there is one small peak at $\sim 1250 \mu\text{Hz}$, reaching

an amplitude of $\sim 7.5 \text{ cm}^2 \text{ s}^{-2}$; however, this peak is no larger than subsequent noise peaks at higher frequencies (1750 μHz and 2100 μHz in particular), and thus, is not high enough above our overall noise floor to be considered significant, nor a detection.

Our results, with a more sensitive instrument, collected on a larger telescope, and with a significantly longer time series for analysis are, therefore, in contradiction to previous results. Therefore, we believe that the excess power previously reported by Gaulme et al. (2011) was perhaps a false detection of the global modes of oscillation, and that the excess power was likely due to some non-Jovian source. If any modes are present within this power spectrum, we believe they are below our final noise floor of 4.5 cm s^{-1} .

4.5 Final PMODE Amplitude Spectra for All Defined Modes of Interest

Next, we will examine the particular modes of interest as selected in Section 3.6.2. These are the modes for which we expect PMODE to be most sensitive, and thus, should be the most likely to reveal global modes, if they are present within these data.

4.5.1 *Final Amplitude Spectra: Low-Order Modes*

First, we display the final, calibrated amplitude spectra for all low-order modes—up to $\ell = 3$ and all corresponding m values—as these are predicted to be the modes where any oscillatory signals will have the highest power. We note that hereafter, we are displaying the amplitude spectra for ease of viewing, as opposed to squaring these to getting the power spectra (as we did in Figure 4.12, for comparison). This subset of modes also includes the amplitude spectrum for the $Y_\ell^m = (1, 0)$, the power spectrum of which has been examined

prior, in Figure 4.12.

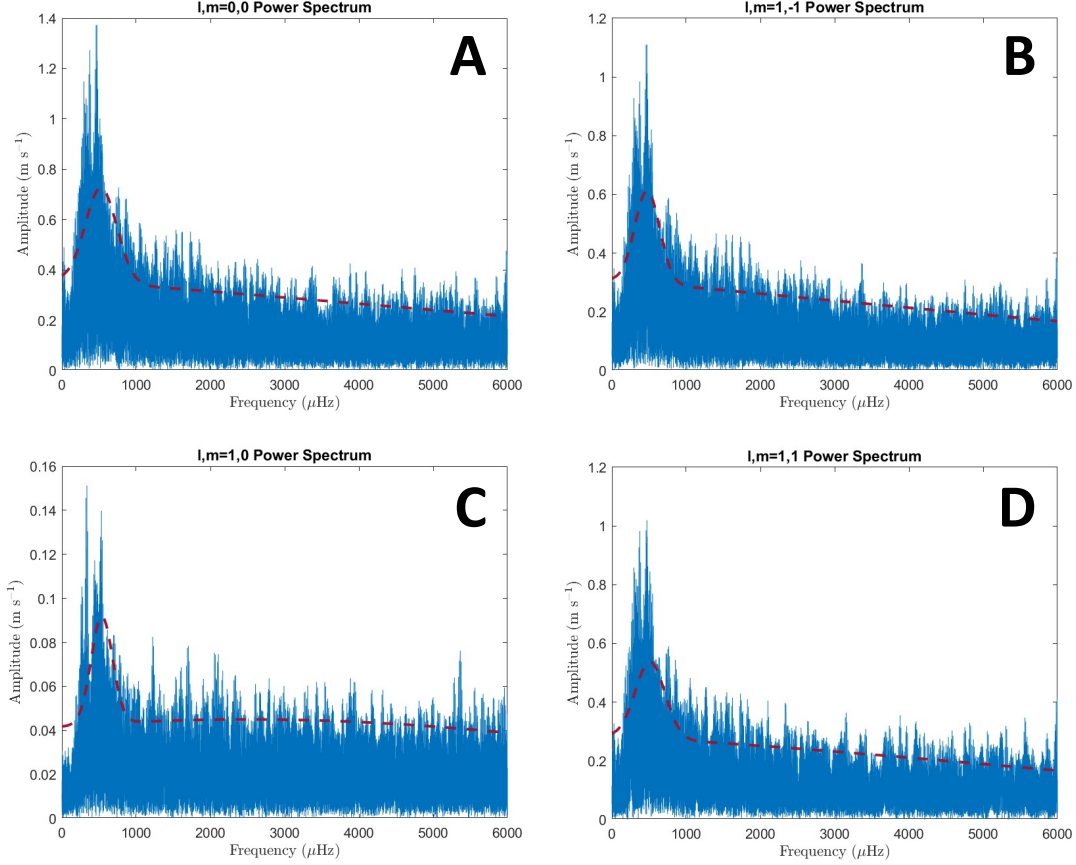


Figure 4.13 Amplitude spectra for all low-order modes ranging from $Y_\ell^m = (0, 0)$ to $Y_\ell^m = (1, 1)$ from the 2020 PMODE observational campaign. The dashed red curve represents our noise floor that varies with frequency as calculated by our “background fitting” technique for each respective spectrum.

These amplitude spectra are shown in Figures 4.13 through 4.16. In each, the background noise level, changing with frequency and calculated by our “background fitting” noise routine, is shown overplotted as a dashed red line. Additionally, in each of these figures, we notice the same lowest frequency (i.e., below 700 μHz) excess power described in Section 4.4, which we attribute to observational, analysis, and upwelling events intrinsic to Jupiter, rather than

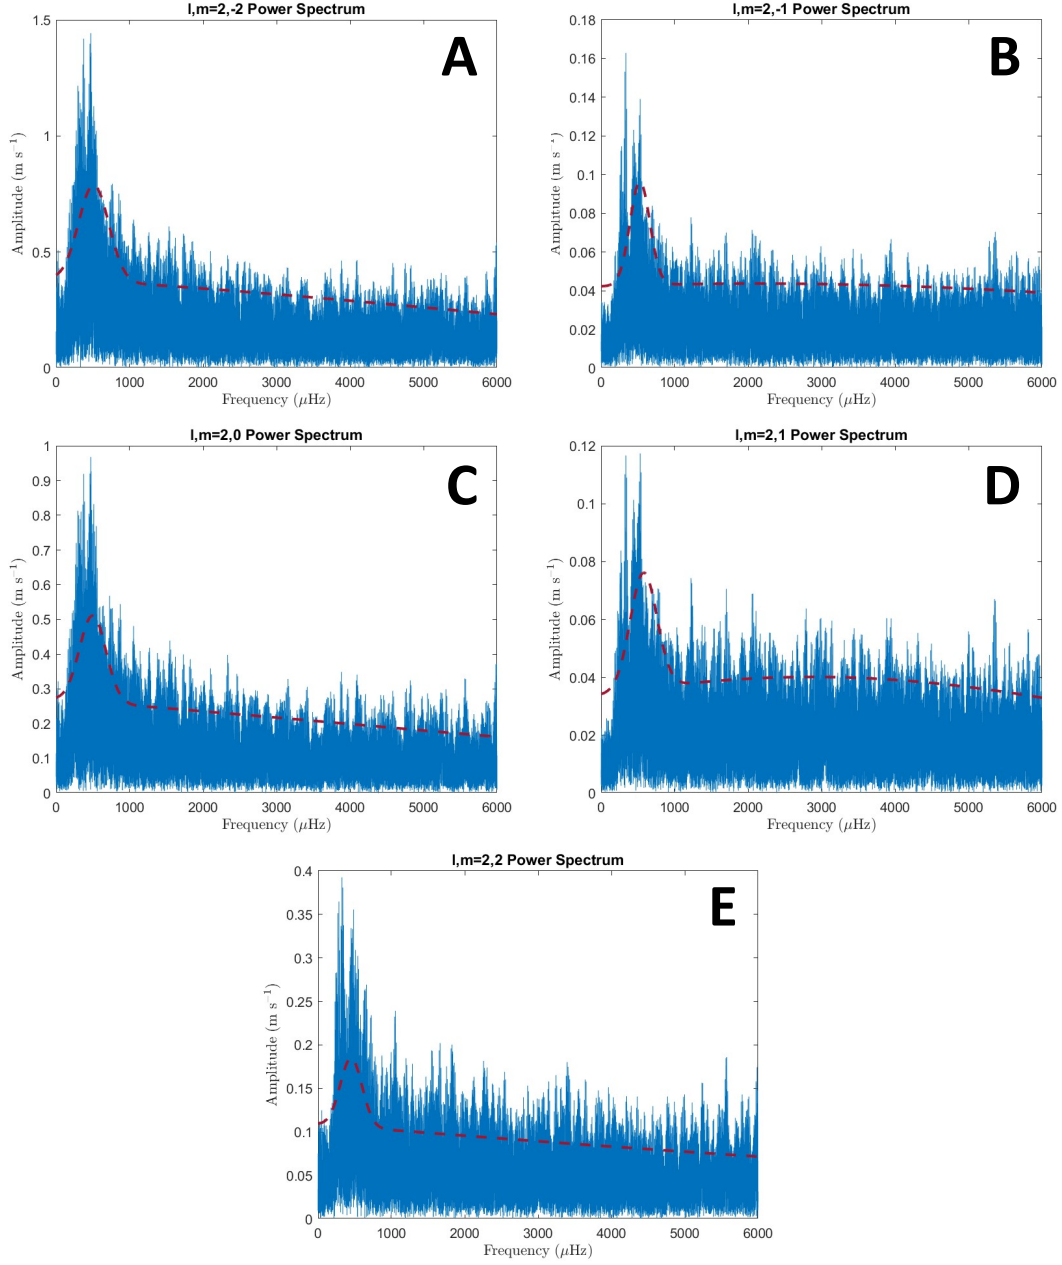


Figure 4.14 Amplitude spectra for all low-order modes ranging from $Y_\ell^m = (2, -2)$ to $Y_\ell^m = (2, 2)$ from the 2020 PMODE observational campaign. The dashed red curve represents our noise floor that varies with frequency as calculated by our “background fitting” technique for each respective spectrum.

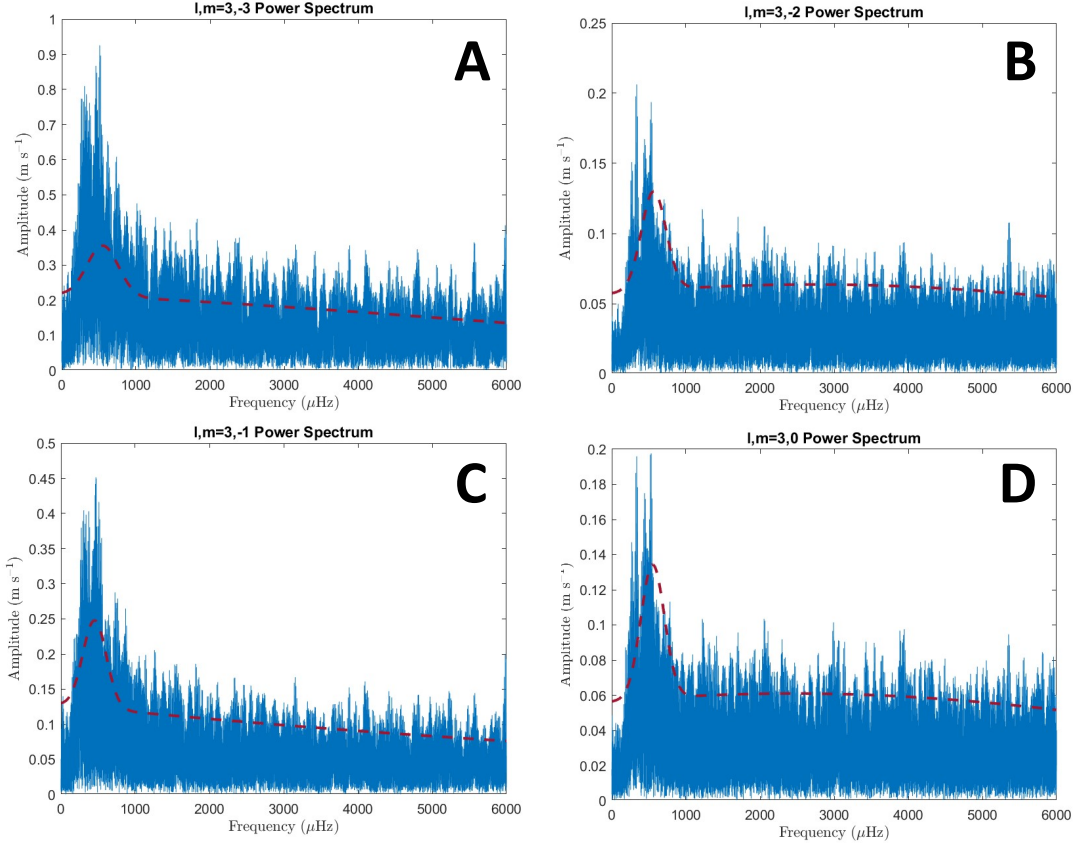


Figure 4.15 Amplitude spectra for all low-order modes ranging from $Y_\ell^m = (3, -3)$ to $Y_\ell^m = (3, 0)$ from the 2020 PMODE observational campaign. The dashed red curve represents our noise floor that varies with frequency as calculated by our “background fitting” technique for each respective spectrum.

to any oscillatory global modes.

We note that the overall noise level reached by PMODE for the $Y_\ell^m = (1, 0)$ mode, shown in Figure 4.13C, is somewhat easier to see than in the comparative power spectrum from Figure 4.12. Here, we see very clearly that there is no significant, organized excess power in the 1200 μHz regime above our noise floor of 4.5 cm s^{-1} , and certainly not above the maximum power within this region, of 8 cm s^{-1} , furthering our confidence that our dataset refutes past claims of modes with amplitudes of up to 50 cm s^{-1} present in this regime.

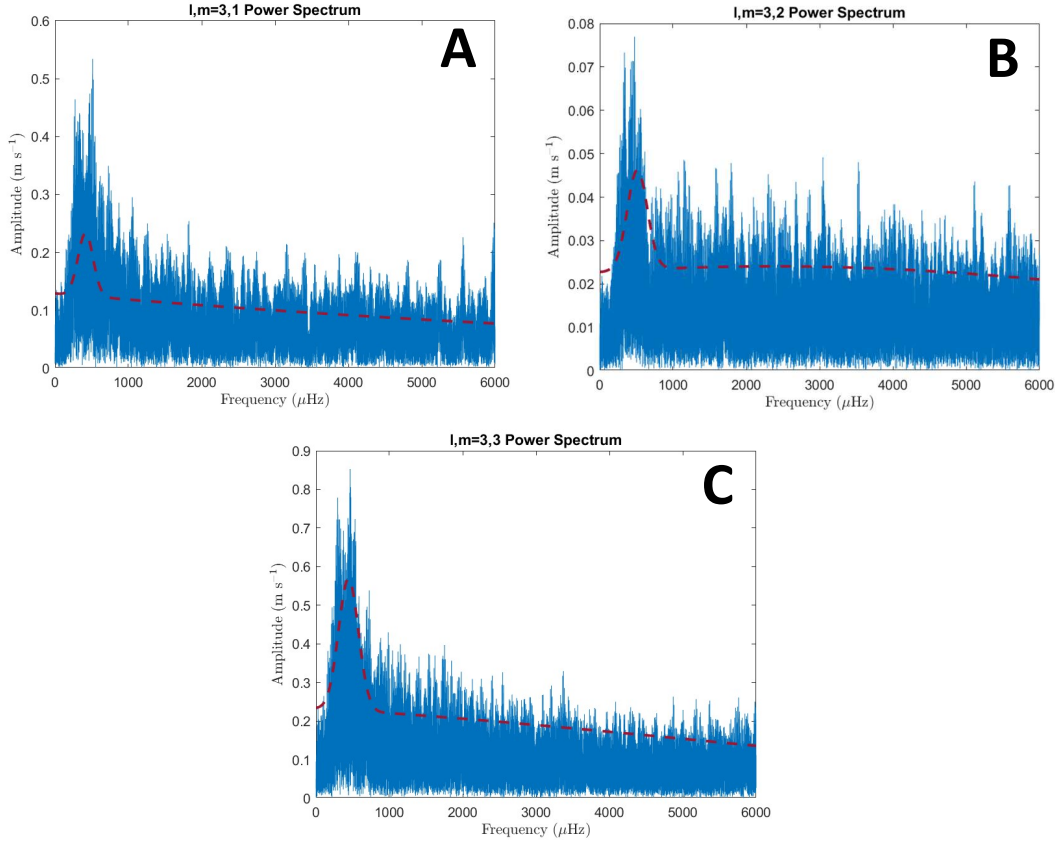


Figure 4.16 Amplitude spectra for all low-order modes ranging from $Y_\ell^m = (3, 1)$ to $Y_\ell^m = (3, 3)$ from the 2020 PMODE observational campaign. The dashed red curve represents our noise floor that varies with frequency as calculated by our “background fitting” technique for each respective spectrum.

Aside from this $Y_\ell^m = (1, 0)$ mode which we have thoroughly analyzed, we note that we see no significant excess power present *anywhere* within any of these low-order spectra. Therefore, we place upper constraints on the maximum possible allowed amplitude for each of these modes to be commensurate with the noise floors reported in Tables 4.1 through 4.4.

4.5.2 Final Amplitude Spectra: $l = m$ Modes

Next, we display the amplitude spectra for all Y_ℓ^m modes where $\ell = \pm m$ for the full dataset, ranging from $m = \ell = \pm 4$ to $m = \ell = \pm 11$, as all values below four are already shown in Section 4.5.2. As a reminder, we examine these modes in particular because our calculated spherical harmonic leakage matrices (recall Figures 3.13 and 3.15) for our spatial coverage across the Jovian disk suggest that these modes should suffer the smallest amount of spatial leakage and subsequent spectral confusion, thus, retaining the most true power in any modes which could be present within the spectra.

This subset of modes also includes the amplitude spectrum for $Y_\ell^m = (11, \pm 11)$, the mode for which PMODE has a unique capability to perhaps be the most sensitive to, as the full amplitude (from receding to protruding portions of the disk as it oscillates) are both contained within the sensitive region of the absorption line. This unique situation implies that these modes may show the full power for any oscillatory mode present, instead of sampling only a portion of its amplitude. As a reminder, a visualization of this effect is shown in Figure 3.14, with further explanation in Section 3.6.2.

These amplitude spectra are shown in Figures 4.17 through 4.19. Again, the background noise level is overplotted as a dashed red line, and the low-frequency excess “noise and upwelling” resultant power are both present. Again, we note that we see no significant excess power present *anywhere* within any amplitude spectra, including the $Y_\ell^m = (11, \pm 11)$ modes of interest. Therefore, we again place upper constraints on the maximum possible allowed amplitude for each of these $\ell = \pm m$ modes to be commensurate with the noise floors

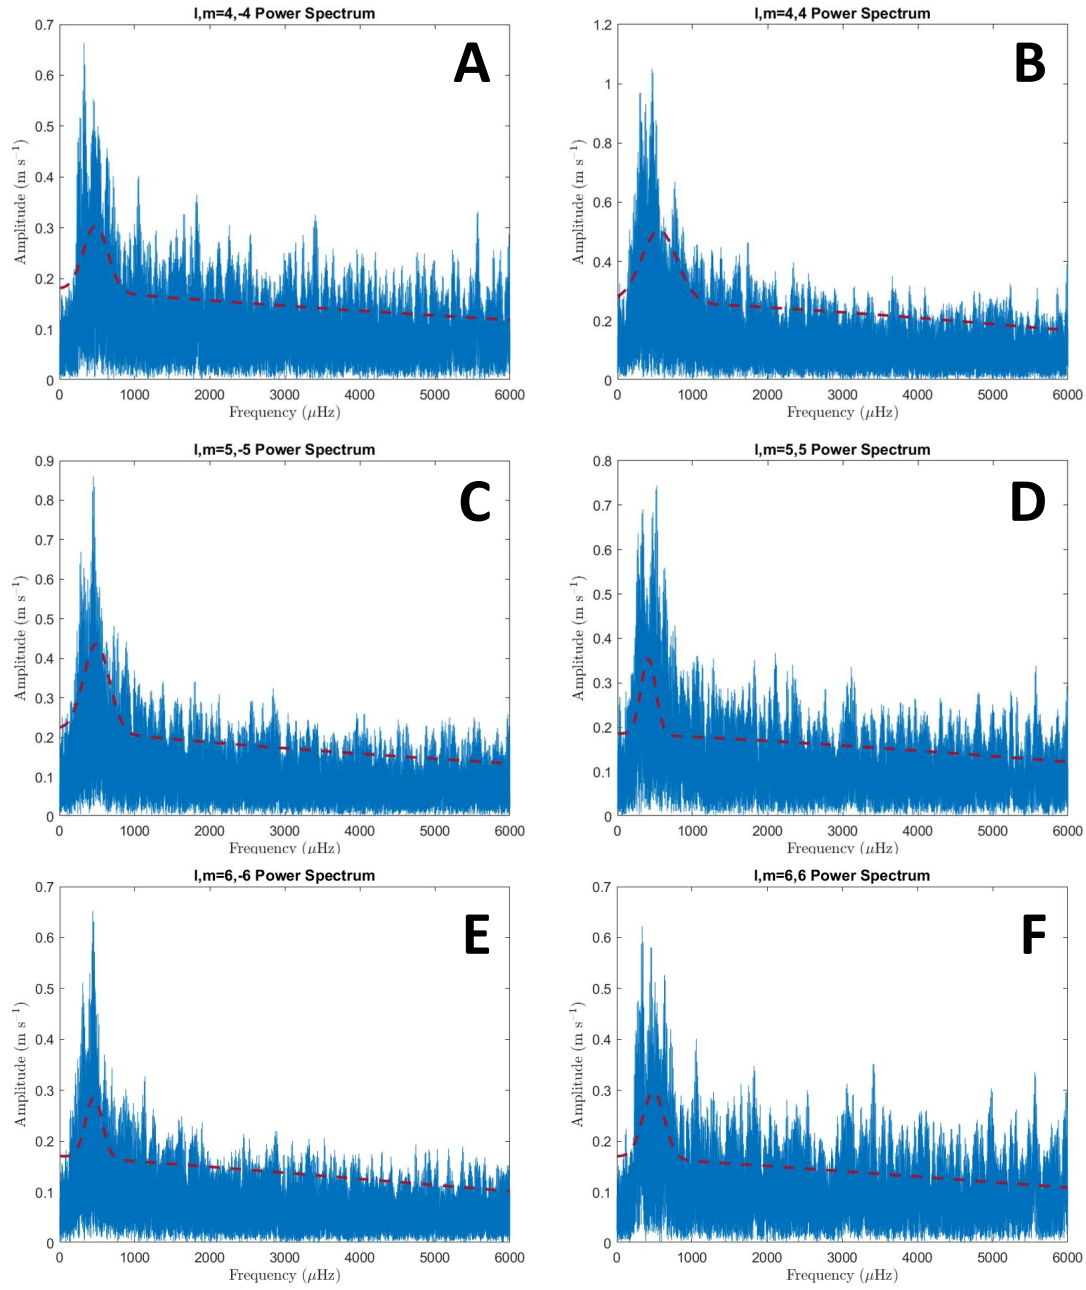


Figure 4.17 Amplitude spectra for the $\ell = m = \pm 4, 5, 6$ modes from the 2020 PMODE observational campaign. The dashed red curve represents our noise floor that varies with frequency as calculated by our “background fitting” technique for each respective spectrum.

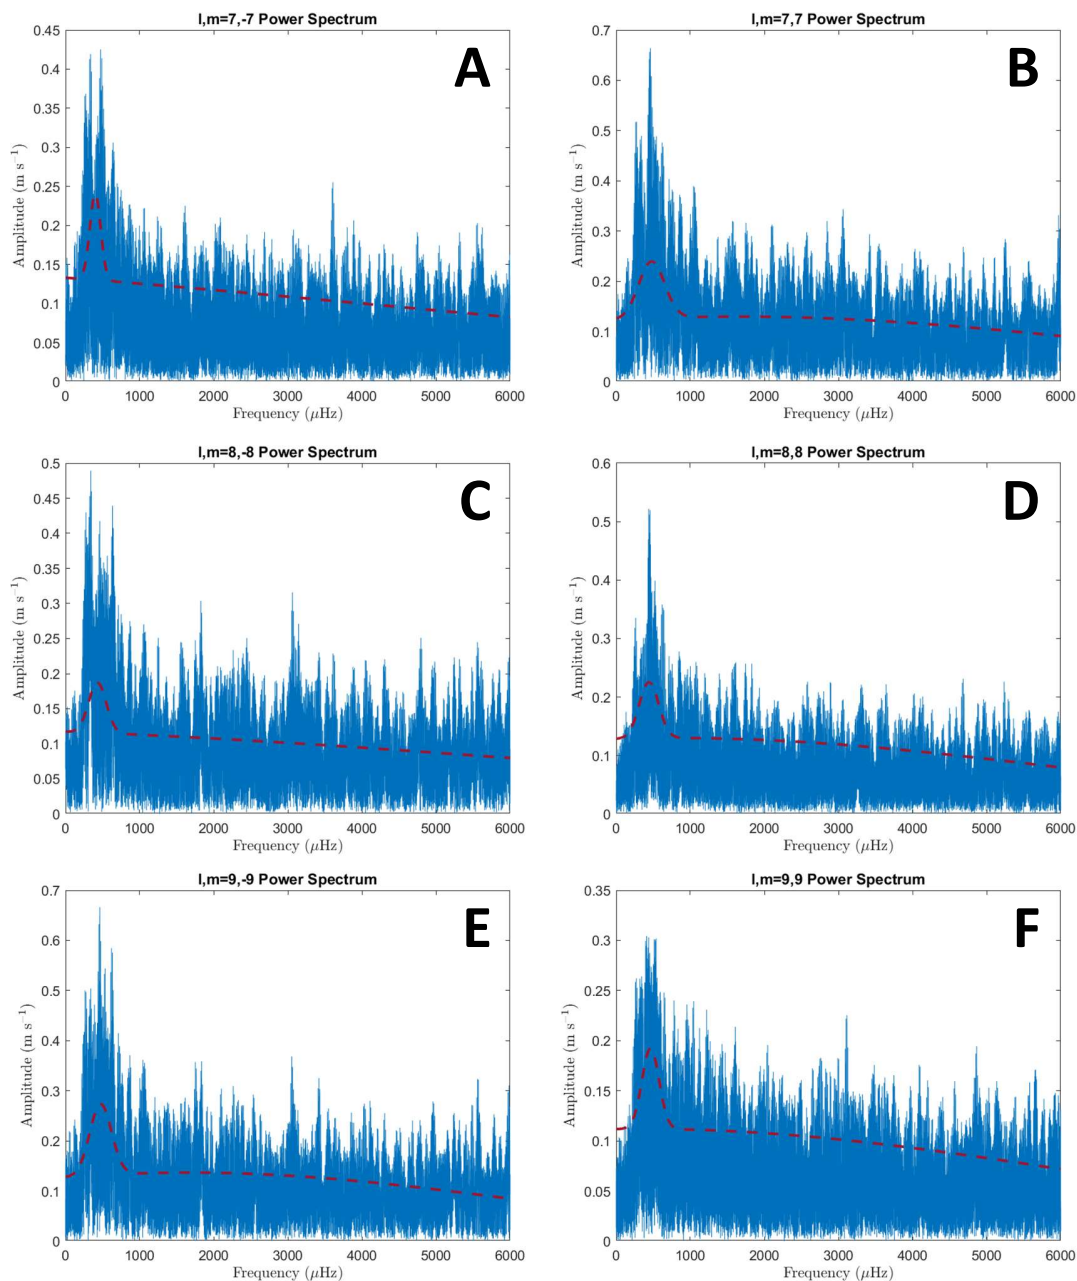


Figure 4.18 Amplitude spectra for the $\ell = m = \pm 7, 8, 9$ modes from the 2020 PMODE observational campaign. The dashed red curve represents our noise floor that varies with frequency as calculated by our “background fitting” technique for each respective spectrum.

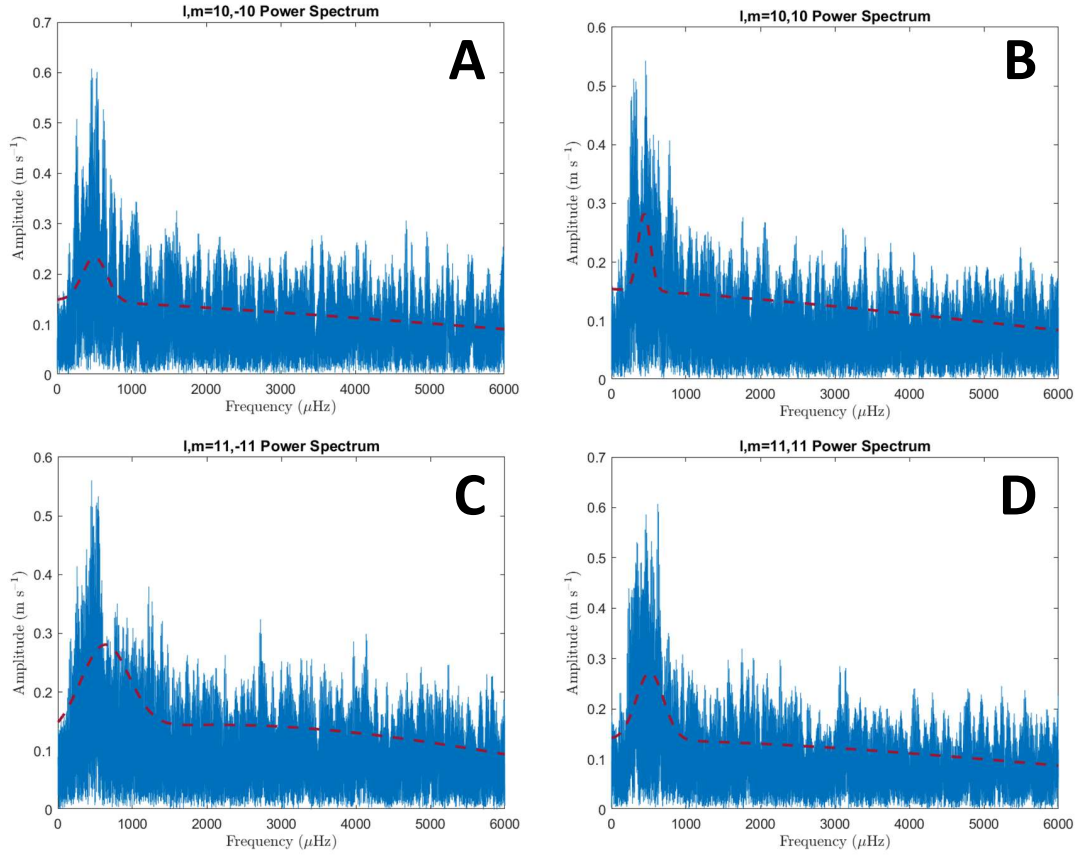


Figure 4.19 Amplitude spectra for the $\ell = m = \pm 10, 11$ modes from the 2020 PMODE observational campaign. The dashed red curve represents our noise floor that varies with frequency as calculated by our “background fitting” technique for each respective spectrum.

reported in Tables 4.1 through 4.4.

4.5.3 Final Amplitude Spectra: Hemispherically Antisymmetric Modes

Finally, we display the amplitude spectra for all hemispherically antisymmetric modes up to and including $\ell = 5$ and all corresponding m values, where the antisymmetric modes are shown circled within Figure 3.9, and italicized in Tables 4.1 through 4.4.

We examine these modes, in particular, because hemispherically antisymmetric modes

result in a net cancellation of any global noise signals, resulting in a substantial decrease in the noise floor for these spectra (this decrease is evidenced by comparing the italicized lines within Tables 4.1 through 4.4 with all non-hemispherically-antisymmetric modes). Thus, these modes should provide an overall lower noise floor, and the capability to detect smaller-amplitude modes, if present within these data.

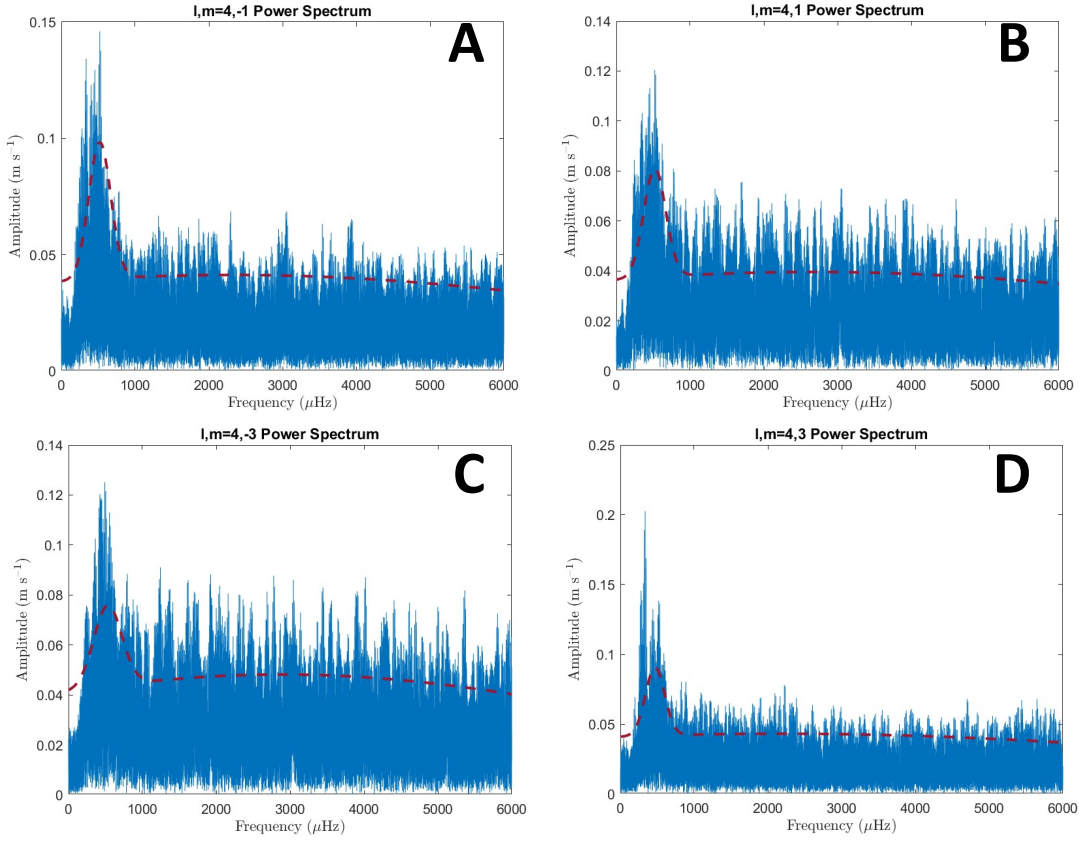


Figure 4.20 Amplitude spectra for the $Y_\ell^m = (4, \pm 1), Y_\ell^m = (4, \pm 3)$ hemispherically antisymmetric modes from the 2020 PMODE observational campaign. The dashed red curve represents our noise floor that varies with frequency as calculated by our “background fitting” technique for each respective spectrum.

These amplitude spectra are shown in Figures 4.20 through 4.21. The background noise

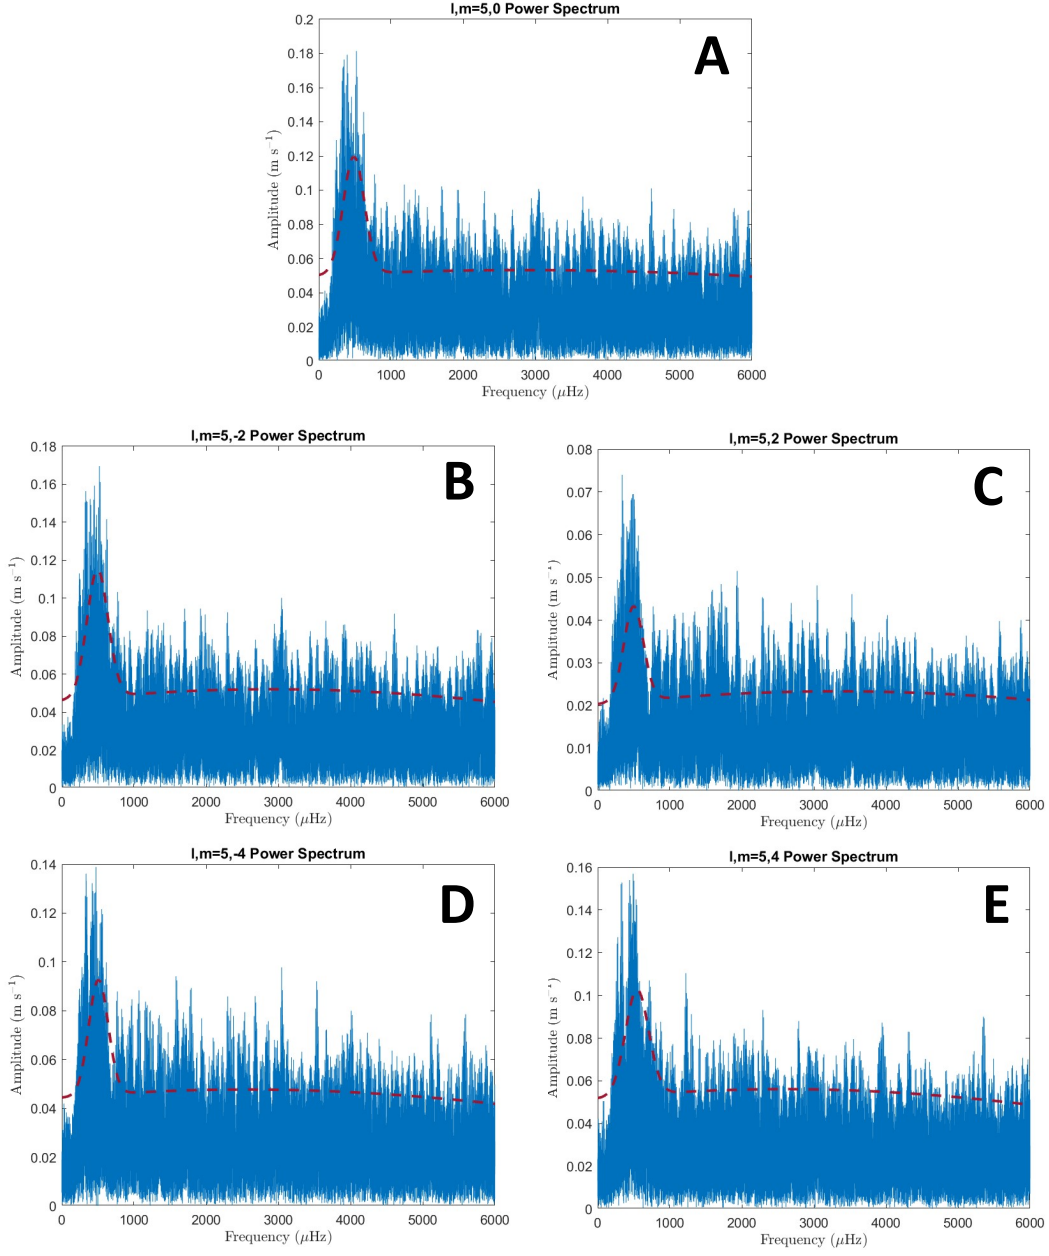


Figure 4.21 Amplitude spectra for the $Y_\ell^m = (5, 0)$, $Y_\ell^m = (5, \pm 2)$, $Y_\ell^m = (5, \pm 4)$ hemispherically antisymmetric modes from the 2020 PMODE observational campaign. The dashed red curve represents our noise floor that varies with frequency as calculated by our “background fitting” technique for each respective spectrum.

level is once more overplotted as a dashed red line, and the low-frequency excess “noise and upwelling” power, as previously discussed, are both present within each amplitude spectrum.

We once again note that we see no significant excess power present *anywhere* within any of the presented amplitude spectra, even for these lower-noise, hemispherically antisymmetric modes of interest. Therefore, we finally, once again, place upper constraints on the maximum possible allowed amplitude for each of these hemispherically antisymmetric modes to be commensurate with the noise floors reported in Tables 4.1 through 4.4.

4.5.4 Review of Presented Amplitude Spectra and Implications for Remaining Modes

All the defined modes of interest that we expect to be most sensitive to have been shown within this chapter. However, these are not the only modes which have been analyzed—we carefully analyzed each of the 144 amplitude spectra, up to $\ell = 11$ with all corresponding m values. We note that neither in the modes of interest explicitly shown in this manuscript, nor the remaining set of modes listed in Tables 4.1 through 4.4 (but without plots shown in this manuscript), do we detect any signal whatsoever that is significantly above our respective noise floors.

As such, we extend our assumption to a final claim that there is no evidence for the detection of global modes contained anywhere within the PMODE dataset. Thus, we report that the noise values reported in Tables 4.1 through 4.4 should hereafter be considered as the maximum possible amplitudes that any global modes ringing within Jupiter may have.

CHAPTER 5

ON PRESENT COMPLICATIONS AND FUTURE PATHS FOR DETECTING THE JOVIAN OSCILLATIONS

Although we have not detected modes with amplitudes greater than our noise floor of 4.5 cm s^{-1} in the Y_ℓ^m mode of interest, we would be remiss to take this as an implication that there are no oscillations propagating throughout the planet. Instead, it is most likely that the oscillations certainly are present, but simply excited by a mechanism depositing lower energy into the modes, such as water storms or turbulent convection.

This is, perhaps, a more easily understandable prediction, as we know both these effects to be actively occurring on Jupiter, as opposed to the previously theorized rock storms, the existence of which has not yet been confirmed on Jupiter. However, the amplitudes for global modes excited by either water storms or turbulent convection are predicted to be on the order of $\sim 0.1 \text{ cm s}^{-1}$, nearly an order of magnitude below what can be best achieved from any single, low-latitude site.

It is clear that, in order to progress current capabilities to detect modes at this level, advancements must be made, in either signal-to-noise in the image domain, or a reduction in spectral confusion in the frequency domain. The latter can be accomplished only by observations from space-based telescopes (the current agreed-upon plan from the *Juno* community), or from polar sites on Earth (Shaw & Gullledge et al., 2023, in prep), both of which will be discussed later within this chapter.

However, developing and launching a space-based telescope, even an Earth-orbiting cube-sat, would take at *least* five years from concept to first light, and significantly longer than

this for a dedicated flagship-level Jovian orbiter. Similarly, Jupiter will not rise above 10° at the South Pole again until the year 2030, and the Northern Arctic circle boasts a lower duty cycle (since observations cannot be conducted from pole itself), worse weather, significantly less infrastructure, and fewer funding opportunities. Therefore, we also consider steps that can be taken in a more timely manner to validate current measurements and improve sensitivity for future measurements.

5.1 Further Validation of PMODE Velocimetric Sensitivity with an Astronomical Scene Generator

During this lull in time before subsequent observations can be collected to progress the field of Diaseismology, we believe it would be valuable to provide a further validation that PMODE is indeed capable of sensing 50 cm s^{-1} global modes, of which we are refuting the previous detection of. Certainly, we have demonstrated PMODE’s sensitivity to velocimetric signals in the clear measurement of the Earth’s rotational signal (Figures 4.1 and 4.2), and our new and sensitive measurement of the Jovian zonal wind profile (Figure 4.4).

However, neither of these detections requires the entire data-processing procedure as thoroughly detailed in Section 3.2. In particular, the average-subtraction, spherical harmonic multiplication, and Lomb-Scargle periodogram analysis, are not applied to either of the data products in which we have clear detections of velocimetric sensitivity. We have applied a significant number of processes to our data, and it is certainly reasonable to have concerns that some of the reduction steps have induced hidden problems or masked overall oscillatory, velocimetric signals. Although we have thoroughly tested the techniques applied and their

ability to preserve true signals, and are confident in our null detection of the global modes of Jupiter, we acknowledge that it would be advantageous to find a way to irrefutably validate our optical path, data collection, reduction, and analysis techniques as a whole to prove that if an oscillatory signal exists in the observed data, it is preserved through all subsequent processing steps.

Achieving this goal requires a source of light that 1.) we know to be oscillating with a well-known frequency and amplitude, 2.) can be propagated throughout our instrumentation, and 3.) is similar enough to our true observations of Jupiter that it can be processed through our developed reduction pipelines without substantial modification. Effectively, we need exactly a Jupiter that is slightly changing in intensity (ideally 50 cm s^{-1}) with a known periodicity (ideally ~ 15 minutes), to reproduce the previously reported detections by Gaulme et al. (2011). Finally, we need the ability to have use of the telescope that the PMODE instrument is installed upon. However, as this is a 4-meter-class telescope, it will be nearly impossible to obtain the on-sky observational time necessary for collecting these desired observations, if such a perfectly oscillating target could be found.

These factors lead us to a novel solution: *simulating* an image of Jupiter, with accurate amplitude and phase, and propagating this light representing the simulated Jupiter throughout our entire optical path and onto our final detectors. This will produce a target, with a fully-known amplitude and phase of oscillatory signals (as these will be pre-defined by the observer simulating Jupiter), that will exactly represent our true-data images, and can be propagated throughout our instrument in the Coudé room of the AEOS telescope without

the need for dedicated, on-sky time.

Therefore, our next goal is to simulate realistic images of any celestial body we wish, which can be degraded with atmospheric effects if desired. We propose to achieve this goal by implementing a technique called Computer-Generated Holography (CGH).

5.1.1 Computer-Generated Holography with a Spatial Light Modulator

The CGH technique we choose to implement makes use of a liquid-crystal Spatial Light Modulator (SLM), a technique detailed in Slinger et al. (2005). CGH is useful, in particular, because no physical object is necessary to generate the optical hologram. Instead, interference patterns are generated which produce the desired hologram, with correct phase and amplitude, following propagation throughout an optical system. Once these patterns are determined, they are then imprinted on the incoming light through use of Liquid Crystal on Silicon (LCoS) SLMs—these modulators are designed to change the wavefront of the incoming light through phase retardance of the incoming wave (Meadowlark Optics 2022).

Incoming light, with a defined polarization state, enters the cover-glass of the SLM, and then into the liquid crystal itself, which has a variable response to applied voltage. When no voltage is applied, the crystal molecules are perpendicular to the direction of propagation for the incoming light, and thus, the largest possible modulation occurs. When the voltage is increased, the liquid crystal molecules tilt until they are parallel with the direction of light propagation, minimizing the resultant phase modulation (see Figure 5.1 for a representation of this effect).

Because each pixel electrode is individually controlled, a variety of resultant phase modu-




Figure 5.1 Liquid Crystal Orientation as a result of applied voltage. The left panel displays the orientation when maximum voltage is applied (resulting in a minimum phase modulation), and the right panel displays the orientation when no voltage is applied (resulting in a maximum phase modulation).

lations can occur, allowing the user to have a good deal in flexibility in resultant modulation. One unavoidable drawback of this design, and in particular, of the discrete electrode “pixels” on the backplane of the SLM, is that there remain gaps between these electrode pixels which reflect unmodulated light back out of the SLM in a bright and undesired spot. This reflection is known as the “zeroth-order spot”, and is typically several orders of magnitude brighter than the desired modulated-light output (Jesacher et al. 2008).

A representation of this effect is shown in Figure 2. In short, the SLM itself works to imprint a desired pattern onto the incident light beam, which manifests as a controlled variation of polarization in the reflected beam. Should this system be constructed correctly, and the desired pattern to be imprinted accurately determined, the user should theoretically

be able to generate an unlimited number of holograms at their desire, which can then be propagated throughout the optical system and into the camera.

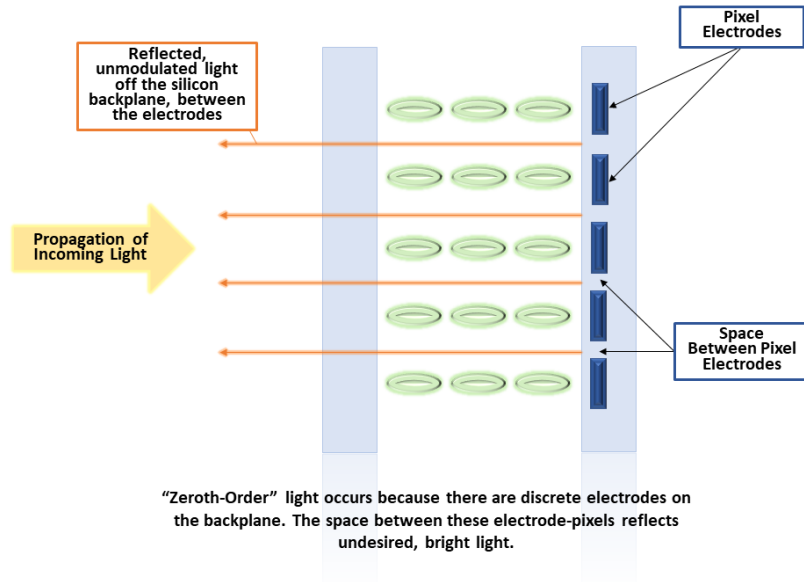


Figure 5.2 The gaps between the “electrode pixels” (blue) cause a portion of the light to be reflected off of the silicon itself (the right-most grey block), resulting in a bright and unwanted “zeroth-order spot”

5.1.2 First Attempt for Determination of Input Patterns for the SLM

Originally, to determine the patterns to apply to the SLM to encode the desired pattern in the reflected light, we chose to use a double pass approach on a single SLM—adding the pattern corresponding to the amplitude of the object on the first pass of light, and the pattern corresponding to the phase of the object on the second pass of light. Naturally, this induces some complexities into the system—primarily, the incident angle of both light beams on the SLM must be less than 10 degrees to prevent incident light impinging on neighboring liquid crystals and their pixels, and subsequently, incorrect modulation of the light.

Secondly, the passes of light on the SLM must be optically conjugate—essentially, the first pass must be re-imaged onto the second pass with no differences in size, focus, etc., to ensure that the phase from the second pass is added at the correct locations to its corresponding amplitude from the first pass. Thirdly, the software must now understand that the SLM is behaving as two separate halves, each functioning independently and as their own entities, instead of together as one large piece to modulate the light.

A significant complication in this task is that of finding a phase and amplitude that will accurately produce the image. For the sake of explanation in this section, we primarily discuss our calibration image, which is a .fits file displaying the HST. Although our end goal is to simulate Jupiter, the sharp edges and clean structure on this image of the HST allows for more careful calibration before transferring to planetary bodies. To obtain the amplitude and phase that produce this image, we follow the following steps:

1. Generate a random array to represent amplitude and another random array to represent phase. Combine the amplitude (A) and the phase (ϕ) to produce the complex pupil (see Figure 2.9 for a refresher on pupil plane vs. image plane) amplitude:

$$U(x, y) = |A(x, y)|e^{i\phi(x, y)} \quad (5.1)$$

2. Autocorrelate this simulated, complex pupil with itself, then take the square modulus of the inverse Fourier transform of this autocorrelation to produce a simulated object image:

$$\hat{g}_k(x) = |FT^{-1}(U(x, y))|^2 \quad (5.2)$$

3. Compare this simulated object image $\hat{g}_k(x)$ with the true object image $g_k(x)$. Minimize the difference between these two images:

$$c = \sum_k \sum_x (g_k(x) - \hat{g}_k(x))^2 \quad (5.3)$$

4. If the difference metric, c , is higher than the desired threshold, generate a different amplitude and phase, and repeat the process.
5. Run an optimizer of your choice (for this experiment, we chose a derivative-based one) to calculate varying phase and amplitude components until the difference between this model image and the input image is minimized.
6. When you are satisfied about the minimal difference between the model phase and amplitude and the input image (approximately 100 iterations for a visually similar image, approximately 1,000 iterations for a physically similar image), these are your input amplitude (Pass #1, hereafter P1) and phase (Pass #2, hereafter P2) for the SLM.

5.1.3 Optical Path for the Astronomical Scene Generator

The original optical path and alignment for this configuration is shown in Figure 4. This system is designed to take in the beam from the AEOS 3.6m telescope on the adaptive optics bench, and subsequently, to output a beam as the same size as the input beam for final transmission through the true optical path that actual observed light through the telescope would see.

To mimic this on the smaller optical bench which we are using for initial alignment, we use a 4-D interferometer with a similar beam size to AEOS. We then add a beam reducing lens to decrease the size of the input beam to fit on the first pass of the SLM (approximately 2mm), then pass this beam through a pupil relay system, then through a first polarizer to ensure all the input light on the SLM has the same polarization, which is 45 degrees with respect to the slow axis of the SLM for amplitude modulation, per the Meadowlark guidelines. (The light then passes through a half-wave-plate, which is only used for calibration in this setup and is not necessary for the optical system.)

Next, we image the pupil on the SLM for the first pass, where we add amplitude, which was calculated via the iterative optimization method described prior via modulation of the SLM. The light then passes through the first of a pair of lenses to create a 4-f relay system—that is, where the light travels through a particular relay arrangement that allows you to produce the final desired output beam size.

The beam then bounces off of two fold mirrors to ensure it is re-directed to the SLM at an incident angle of less than 10 degrees. The light then passes through a second polarizer, crossed relative to the first, to block any light that has not been modulated by the SLM—thus, only the amplitude-modulated light should be allowed to pass. Importantly, this should block the “zeroth-order aberration” (recall Figure 5.2).

Next, the beam passes through a second half-wave-plate, this time necessary for the path. This second half wave plate rotates the output polarization to be parallel with the slow axis of the SLM for addition of the phase component, which again was calculated via the iterative

approach described earlier. The beam then passes through the second portion of the 4-f relay to re-image the first pass of the SLM onto the second pass, then bounces off the second pass itself on the SLM, where the phase component is added. After this second pass on the SLM, the light passes off of two more fold mirrors, and then into the imaging lens. The optical path for this setup is shown in Figure 5.3.




Figure 5.3 The original optical path for the SLM object simulator on the optical bench.

Output images from this setup are shown in Figure 5.4. Panel A displays a reconstruction of the HST from the input phase and amplitude, and Panel B displays the Fast Fourier Transform (FFT) of a star, from the image of a star itself being generated on the phase side of the SLM (as opposed to the star generated from its corresponding phase). These images display some capability of this described setup to produce output results, although they are

not of the high quality as expected. Primarily, the image of the HST displays a large, overall speckled appearance, which is not present within the input image itself.




Figure 5.4 Output images from the current setup for the SLM, as shown in Figure 5.3. Panel A displays a reconstruction of *Hubble* from the input phase and amplitude, and Panel B displays the FFT of a star, from the image of a star itself being generated on the phase side of the SLM (as opposed to the star generated from its corresponding phase). These images display some capability of our current setup to produce output results, although they are not of the final quality as expected.

5.1.4 Initial Complications and Second Attempt for Determination of Input Patterns for the SLM

Further investigation into the current state of the field reveals that this unexpected and undesired speckled appearance in our output images is in fact not an artefact due to improper alignment or input SLM pass generation, but instead, an unavoidable addition of noise added by this previously detailed approach—specifically, the addition of an amplitude directly

onto the SLM, which is a phase-modulation device, induces an additional “noise phase” into the system that is unaccounted for elsewhere. This results in an increased overall speckled pattern, and additionally, a very high contrast between the zeroth-order spot and the simulated background image, resulting in a loss of control over contrast and a sacrifice in the majority of the intensity of the true image (Jesacher et al. 2008).

To mitigate this, we can implement a different method, where instead of applying an amplitude pass and a phase pass directly to the SLM, we apply two phase passes instead. In particular, we run a second solve to calculate a phase which will reproduce the desired amplitude for the first pass (hereafter, we will refer to this “amplitude-phase” as ψ). Additionally, we will also solve for the undesired “noise phase” (hereafter referred to as ϕ) that will be induced from this first pass, which we can subsequently subtract out of the latter, true phase pass (hereafter referred to as θ). Therefore, our new P1 for the SLM will be only ψ , and our second, P2, will now be $\theta - \phi$; we will hereafter refer to this quantity as β . Importantly, P2 on the SLM must now be applied at a focus rather than at a pupil, since these two phases must be applied at optically conjugate Fourier planes (Jesacher et al. 2008).

This second solve is conducted via the exact same aforementioned minimization routine—by generating a random phase and amplitude, producing a complex pupil from these components, and determining the output, model image (however now, your output, model image is the input amplitude $|A|$ from the initial solve. This time, we solve with a constant amplitude, since we want to reproduce the input using phase only (in addition to retrieving the second induced noisy phase θ alongside our desired “amplitude-phase” ψ). We can invoke

prior knowledge that the output amplitude this time must be real and positive. To do so, our input equations to be minimized are as follows:

1. Minimize the difference between the input image, $|A|$, from the model image, $|A_{mod}|$
2. Here,

$$\hat{A}_n = e^{-i\phi_n} \sum_m \beta_m e^{+i2\pi nm/N} \quad (5.4)$$

where ϕ is our “incuded noise phase” as described earlier, i represents the imaginary number, N represents the number of pixels on one side of our input image, and n and m are spatial coordinates.

3. Here, $\beta_m = |B_m|e^{i\psi_m}$, where B_m represents a spatially constant amplitude as described earlier (this is important, because we want solely the phase that will represent the amplitude), and ψ represents our desired phase to accurately reproduce the amplitude.
4. Once quantity 1.) is minimized, we will then have our two output phases— ψ and ϕ —with which we can generate our two passes for the SLM.

In theoretical restorations, this procedure produces a clean and accurate simulation which appears to be a good representation of the input image, although the optimization takes a significant amount of time to run. To minimize this amount of time, we can combine both solves—the initial solve to find the phase and amplitude of the input image, and the secondary solve to find the phases which represent the determined amplitude—into one final set of equations to be minimized simultaneously. We know that the value of one single pixel

in an object (O) on a grid of $N \times N$ pixels can be represented by:

$$o_n - a_n a_n^* \quad (5.5)$$

where $*$ denotes the complex conjugate, and

$$a_n = \sum_m \alpha_m e^{+i\beta_m} e^{+i2\pi mn/N} \quad (5.6)$$

Here, β represents the addition of the true original phase, θ , and the undesired but inevitably induced “noise phase” ϕ , the positive in the second exponent restricts the values from including negative or zero values, and l , m , and n represent 2-D spatial coordinates.

Finally:

$$\alpha_m = \sum_l |B_l| e^{i\psi_l} e^{+i2\pi lm/N} \quad (5.7)$$

Here, $|B_l|$ represents constant amplitudes over the range of the input image, and ψ represents the phase that accurately represents the amplitude. Once we have determined our value for a_n , we subsequently know a_n^* , and thus can minimize our final error metric of:

$$\epsilon = \sum_n (o_n - a_n a_n^*)^2 \quad (5.8)$$

This minimization will provide us with two different phases: P1 containing ψ (the phase representing the amplitude of the input image, and applied to the first pass), and P2, containing β (the combination of the true original phase, θ , and the subtracted “noise phase” ϕ induced from the initial pass). These outputs are shown in Figure 6.

Although this process produces exquisite theoretical reconstructions, we are unfortunately unable to reproduce these results on the optical bench. In particular, the first pass

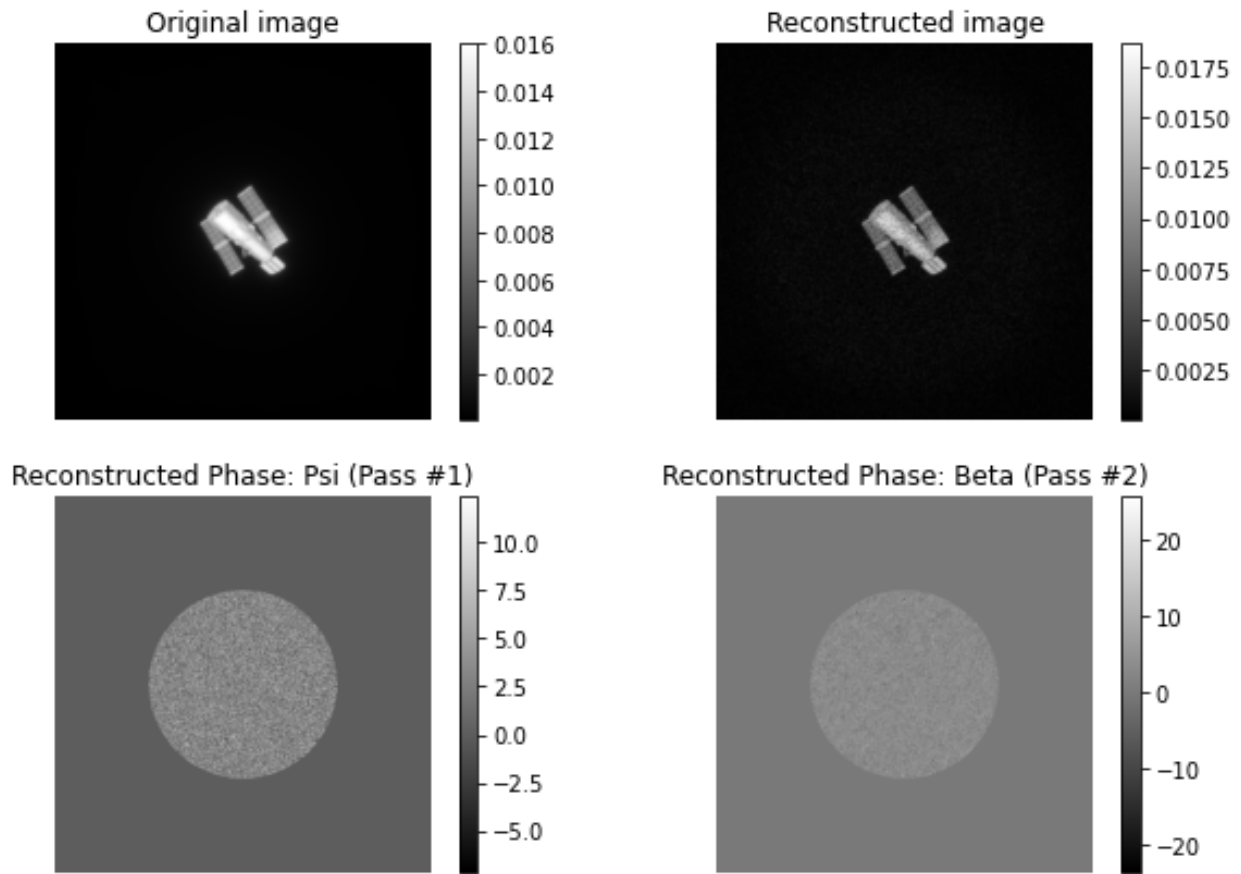


Figure 5.5 Top Left—Input image into the minimization routine. Bottom Left—the first pass on the SLM, containing amplitude information as a phase. Bottom Right: second pass on the SLM, containing phase information. Top right: Reconstructed Image from P1 & P2

(in this case, the FFT of the amplitude is placed onto the first pass, which gets Fourier-transformed as it passes through a lens, producing the amplitude itself falling directly onto the second pass) propagates exactly as desired, reconstructing the amplitude beautifully onto the second pass on the SLM. However, we have currently been unsuccessful at achieving an accurate combination of the two paths—the combination of light thus far entirely scrambles the image, with no resemblance of the input image reaching the detector itself.

We have attempted a variety of techniques to alleviate this complication—full rearrange-

ments of optical systems, replacement of elements which could be causing the fundamental source of the problem, applying restricting irises into the optical path to eliminate any external high-frequency noise that could be destroying the signal, and trying many different phase-retrieval techniques, from the original Bochner decomposition (Abbott 2022, in prep), to a classic Gerchberg-Saxon phase retrieval algorithm (Jesacher et al. 2008), to implementing Artificial Intelligence / Machine Learning (AI/ML) techniques with PyTorch (Wu et al. 2021). However, through all these attempts, we have yet been unsuccessful.

5.1.5 Third Attempt for Determination of Input Patterns for the SLM

As all aforementioned techniques have yet proven unsuccessful, we instead return to a simplified technique: implementing one single pass on the SLM, and applying only phase-modulation to reproduce the desired output image. To do so, we apply the original Bochner solve as detailed above, in Section 5.1.2. However, now, we place external constraints that the amplitude we are solving for is flat, or a circular aperture with values of 1 inside this circular pupil, and values of 0 surrounding it. This has the effect of forcing the solve to incorporate as much information as possible in the phase-only solution to restore the original image. An output image using this technique is shown in Figure 5.6.

Although we note that there are certainly still some speckles present in this simulated image of *Hubble*, we believe this is a drastic improvement over the original output, shown in Figure 5.4, and small details are beginning to appear on the Solar panels and the tail of the satellite. An example of simulating Jupiter with this same, phase-only technique is shown in Figure 5.7.




Figure 5.6 A simulated hologram of the HST, created from applying a phase-only with constant amplitude solution onto the SLM from our original Bochner decomposition. The zeroth-order spot is present in the center of the satellite. This image has been blurred with a 1-sigma Gaussian to alleviate some speckling. Although some speckles are still present, the simulated output should be of sufficient quality for future testing of PMODE.

This reconstruction is not perfect; however, we believe it will be sufficient for installation on the telescope to allow for a future, holographic-target, simulated time series to be propagated directly through the PMODE optical path, providing a new avenue to validate the instrument’s sensitivity and capabilities, as well as establishing even further confidence in our null detection of the Jovian modes of oscillation.

5.2 Future Instrumentation Improvements

In addition to irrefutably validating our instrumentation’s capability to sense oscillatory signals through use of this benchtop simulator, it is also useful to begin looking towards modifications which can be applied to both current and future instrumentation for further investigation of Jovian seismology.




Figure 5.7 A simulated hologram of Jupiter, created from applying a phase-only with constant amplitude solution onto the SLM from our original Bochner decomposition. The zeroth-order spot is present in the center of the satellite. This image has been blurred with a 1-sigma Gaussian to alleviate some speckling. Although some speckles are still present, the simulated output should be of sufficient quality for future testing of PMODE.

5.2.1 Increasing Spatial Coverage on the Jovian Disk

First, a simple modification to improve our signal is through addition of a second velocimetric channel to expand spatial coverage. This would substantially reduce spatial leakage in our power spectra, ensuring that any power present in oscillatory modes remains in its true location. Another benefit of adding a second, broader line, is that its corresponding absorption profile will be significantly easier to disentangle from any average Dopplergrams to reveal the Jovian ZWP, perhaps allowing for a direct measure of the profile without the need for any high-pass filtering.

Therefore, a measurement such as this should very closely reproduce the profile obtained from cloud tracking measurements. As a reminder, this was attempted in our PMODE

campaign (Section 2.2.2), but ultimately, was unsuccessful due to detector artefacts (Section 2.4). Ensuring that both paths were collected successfully with no complicating time- and spatially-varying patterns would be a high priority to increase spatial coverage if awarded further observational time.

5.2.2 A Novel Imaging Shack-Hartmann Wavefront Sensor

Secondly, one of the best ways to improve S/N in astronomical observations is by adding wavefront correction before the aberrated light is incident on the detector, thus ensuring all photons are in their proper location and the light is not spread out by the atmospheric turbulence it travels through. However, adding adaptive optics to conduct this wavefront correction in real time is not only expensive, but it also requires many photons from a very bright point source, an expensive and fragile deformable mirror, a wavefront sensor to measure the aberration, and a fast readout on your detector to correct the light before the atmosphere changes.

Not only are many of the aforementioned stipulations expensive, but one—requiring a point source in the field—is simply not possible when observing an extended target such as Jupiter. In this case, your Point-Spread Function (PSF) will change over the field. In order to accurately measure this PSF over an extended source, another type of wavefront measurement must be collected—expanding the typical “spot diagram” Shack-Hartmann Wavefront Sensor (SHWFS) (Platt & Shack 2001) to magnify each of the individual small spots to an image of the target itself with greater than 30 pixels spanning its diameter, as opposed for the typical 2-3 pixels for a traditional SHWFS.

This magnification of the spots produces a different instrument—now an Imaging Shack-Hartmann Wavefront Sensor (ISHWFS) (Li et al. 2021). A more complex numerical solve can then be implemented to recreate the wavefront over the extent of the observed source. However, this solve requires time, and thus makes it less preferable for real-time adaptive optics correction. For this reason, we turn to look at image restoration in the post-processing domain. If these wavefronts can then be measured and recorded simultaneously to image capture, the PSF for each image can later then be restored, and subsequently de-convolved from the image itself, significantly improving overall S/N.

For the primary historical attempts to measure the Jovian oscillations (this work, as described in Gulledge & Shaw (2022), and the work by Gaulme et al. (2011)), the current adaptive optics and wavefront sensing techniques are limited to only tip/tilt correction at best, and occasionally limited to no real-time wavefront correction at all, depending on observatory capabilities. To this end, we believe that implementation of an ISHWFS to be used for future wavefront restoration of Jovian data will provide a substantial benefit.

Our belief is that this post-processing technique will not only improve any measurements taken from the ground in the next 5–10 years, but will also provide a significant and necessary improvement for future polar sites for Dionesimology, as Jupiter reaches a maximum amplitude of $\sim 22^\circ$ from the South Pole, and spends the remainder of its time above the horizon at a lower altitude, and thus, with light propagating through a substantial amount of atmosphere. Therefore, a functional and concept-proven ISHWFS for Jupiter will be invaluable, and current instrumental designs have a space on the breadboard pre-allotted for

just such an instrument (Shaw & Gullledge et al. 2023, in prep).

Finally, the aforementioned and already-developed SLM from Section 5.1 will be valuable in future development, testing, and proof-of-concept for an ISHWFS. Similarly to our plans for simulating an oscillating Jupiter and propagating its light throughout our optical path, the implementation of this SLM also provides us with the ability to apply atmospheric turbulence effects to these simulated Jupiter holograms at any level we want. We can then directly install a future ISHWFS on the breadboard and send a portion of this simulated, aberrated light into the wavefront sensor to prove how well it is capable of correcting aberrations on an extended target, where the observer has full knowledge of the input aberrations, but the ISHWFS and subsequent reduction software does not. This will provide the capability to validate that we will be able improve our future S/N from Polar sites when Jupiter suffers from significant atmospheric degradation.

5.3 Future Paths for Detecting the Jovian Global Modes of Oscillation

Aside from current steps that can be taken in the immediate future before more dedicated Jovian seismology observations can be collected, it is also worthwhile to look towards the future and begin planning for these next measurements. If the goal of detecting seismic signals within Jovian data was unobtainable by a 24-night, consecutive observational campaign on a 3.6 meter telescope, it is unlikely that reaching this goal will ever be achievable from a single, low-latitude site.

The truth of this becomes even more apparent when we consider the fact that we now

expect any oscillations present to have maximum amplitudes of $\sim 0.1 \text{ cm s}^{-1}$ per current modeling techniques, nearly an order of magnitude smaller than the capabilities achieved *via* PMODE. We can estimate the length of observational time required to reduce our sensitivity with the current instrumentation by a factor of 40, which would be required to attain a noise floor on the order of $\sim 0.1 \text{ cm s}^{-1}$. We know that signal scales as time squared—meaning that we would require $S_2 = S_1 \times t^2 = 24 \times 40^2 = 38,400$ full days of observational time, or approximately 105 *years* of observations on a 3.6m telescope. Naturally, this is entirely unfeasible.

Instead, we look to alternative and exotic solutions to reduce our noise floor for future observational campaigns. One way of achieving this is by reducing the spectral confusion induced into our dataset through avoiding periodic gaps caused by the rise of the Sun each day. This complicating factor, although inevitable in observations from a low-latitude site, can be avoided by placing a telescope at one of the poles of the Earth. Historically, collecting time-series observations from the South Pole for this exact reason of avoiding spectral confusion has resulted in the most significant breakthroughs for the field of Helioseismology; as such, it is safe to assume that Polar observations would provide the best Earth-based chance at detecting Jovian oscillations, as well.

Secondly, a clear way to avoid daytime gaps in data collection is to observe your target from space. A satellite orbiting around Jupiter itself, placed in the second LaGrange point, or even orbiting the Earth in a Sun-synchronous orbit would all provide near-fully-continuous observations of Jupiter, drastically reducing the noise floor and spectral confusion better

than even Earth-based Polar observation sites. A space-based solution such as this has the added benefit of no complicating atmospheric noise, greatly increased resolution, less overall light-loss, and would have significantly fewer optical components than many Earth-based observatories; thus, much less light would be lost simply to optical transmission.

These possibilities are discussed in detail within the following sections.

5.3.1 Adaptation of the PMODE Instrumentation for Deployment to the Geographic South Pole

PMODE evolved from an original design, which was created with the intent of an eventual deployment of the prime instrumentation to the geographic South Pole. This site will allow for approximately 100 full days observing with gaps only due to irregular weather sources, eliminating the periodic daytime gaps which we know to inject deleterious side-lobes into the time series. Although the instrumentation and optical paths are effectively the same, an observation site such as the South Pole is a much harsher environment than the Coudé room of a weather- and temperature-controlled telescope.

Achieving this goal requires not only design, alignment, and build of an optical instrument, but of an entire *observatory* itself, with a full telescope, mounting for the optical breadboard, computers for telescope control, a self-sufficient power supply, temperature regulation, extremely effective insulation, and plenty of back-ups should anything fail. However, the benefit and payoff of transitioning to a Polar site is unparalleled by any other observations achievable from the ground.

Although this is a lofty goal, it is not *just* a goal—this observatory has not only been

conceived in idea; rather, it has been conceptualized, constructed, aligned, and tested, and is ready for deployment to a Polar site at the first available opportunity. This Polar version of our PMODE instrument has been christened as the Long-duration Antarctic Night and Day Imaging Telescope (LANDIT). LANDIT boasts the same PMODE instrumentation, albeit mounted on a much smaller-footprint breadboard, which is subsequently located on the Cassegrain port of a 0.5m Ritchey–Chrétien telescope, all controlled upon a Software-Bisque Taurus 700 Alt-Az telescope mount.

The small-footprint breadboard for LANDIT has already been shown, in Figure 2.7. The full observatory, in all her glory, is shown in Figure 5.8. In the second panel of this image, we see the breadboard installed directly on the backplane of the telescope itself. The telescope is encased in a lightweight, carbon-fiber protective enclosure (better known as the “Hex Tube”), which is lined within with blue foam to act as an insulator (a necessity, given the sub-zero temperatures at the South Pole and our Doppler velocimeter’s absolute requirement for stable temperatures).

Also encased within this Hex Tube are several heaters, two of which can be seen embedded in the cords in the rightmost panel of Figure 5.8. This heating system is completed through addition of temperature probes, Proportional Integral Derivative (PID) controllers to test and maintain the desired temperature, and additional mounted fans to keep the warm air circulating throughout the interior of the optical tube. A counterweight system is mounted onto the metal interface plates on the sides of the Hex Tube, which connect the fork-arms to the telescope motors.





Figure 5.8 LANDIT assembled—here, we see the 0.5m telescope, complete with the blue-foam and carbon fiber insulating outer structure. The dual-paned entrance aperture is removed in these images. On the right, we are observing the telescope from behind, where the optical breadboard is shown mounted to the back-plane of the telescope; here, the enclosing rear carbon-fiber panel has been removed to show the interior. Light enters out of the page through a central hole in the breadboard, after which it is relayed upward through the scientific instrumentation that is displayed within Figure 2.7. All electrical wiring passes through a bulkhead in the lower right portion of the Hex Tube; in the final installation, the wiring bundle will be wrapped a soft insulating case until it is reconnect through a second bulkhead in the insulated fiberglass base, and finally connected to the enclosed power source.

In this figure, the front panel is missing the glass window(s) acting as an entrance aperture for the telescope. This window, with no additional features, would certainly act as a heat sink for the telescope, and much of the internal heat of the Hex Tube would be lost through it. To solve this, we instead opt for a dual-window system with a gap between both glass panels—this gap is then filled with an inert gas (nitrogen, for our purposes) to act as an insulator while still allowing all desired light to pass through. Finally, on the front of the telescope, the small finder-scope window is shown directly above the entrance aperture; this window is also dual-paned and filled with inert nitrogen to prevent heat loss as well.

This upper hex tube is encased with a final carbon fiber back-panel, enclosing the instrumentation on the rear of the telescope, and all wiring is led through a bulkhead mounted on the side, through an insulating cable wrap, and eventually, through a second bulkhead and into the base enclosure of the telescope. This base enclosure is made of fiberglass (from a local sailboat manufacture on Maui Island) as opposed to carbon fiber, as a lightweight enclosure was still necessary for the overall structure; however, the base does not have as strict of weight requirements as the Hex Tube. As with the upper tube, this base is lined with insulating blue foam as well, and houses the guidance and storage computers for the telescope, as well as heater controllers and connections for power cabling. The inclusion of all these factors produces a lightweight, deployable, complete, and unique observatory, capable of withstanding the extreme weather at the geographic South Pole.

This system has been fully assembled, attained first light images, and has passed rigorous tests to ensure that it meets all defined thermal requirements. There are minor issues within

this observatory which remain to be fully mitigated prior to deployment—in particular, the performance of the Software Bisque mount was less than ideal and resulted in significant overall jitter of the telescope; further assessment and testing of vibrational damping to mitigate this complication are likely required before this telescope will reach the Pole. However, once funding is attained, this self-sufficient observatory is fully constructed and can be prepared for deployment to an Antarctic or Arctic site with only a few months required for final preparation and testing.

5.3.2 LANDIT Thermal Testing

As the exceptionally harsh environment of the South Pole reach average minima of -80° F (or -62° C; values obtained from historical National Oceanic and Atmospheric Administration (NOAA) data) in the Austral winter (when LANDIT would be observing), it is necessary to validate that our instrumentation is capable of maintaining a temperature stability of $\pm 1^{\circ}$ C when subjected to a large temperature difference to ensure that the optical components do not fail and that the MOF heaters surrounding the vapor cells are not subjected to an excessive amount of stress to maintain the exceptionally stable gas temperatures for seismic measurements. However, the island of Maui—where we have constructed, assembled, and tested the LANDIT observatory over the past four years—is not home to the extremely cold temperatures necessary for testing.

To find the coldest possible site for LANDIT, we contacted a local seafood restaurant—Eskimo Candy—which is home to a industrial size and strength deep-freezer. As the Hex Tube itself (and the optical instrumentation within it) is the the most sensitive to tem-

perature swings, we isolated this component in particular, removed the telescope structure and the optical breadboard for safety, and ensured the full, self-sufficient heating subsystem (including thermometers, PID controllers, heaters, and fans) was running correctly. We connected a laptop computer to the heating system, sealed all these components within the insulated Hex Tube, including the frontal window panel and rear carbon-fiber panel, placed the entire assembly on a shelf adjacent to a few slabs of frozen Ahi, and left the system to run overnight.

Even in this deep freezer, the minimum attainable temperature is -15°C . To create a temperature difference more reminiscent of that which we would be subjected to at the South Pole, the heating system inside the sealed Hex Tube was directed to maintain an internal temperature of $+30^{\circ}\text{C}$, as opposed to the defined $+3^{\circ}\text{C}$ requirement for the final South Pole system. This produces an overall delta of 45°C —not as extreme as the expected 83°C that our instrumentation must withstand at the South Pole, but providing the best means of validation currently achievable from the Hawaiian islands.

The results from this thermal testing are shown in Figure 5.9. Here, the black line represents the measured temperature in the back portion of the Hex Tube, and the surrounding grey regime represents the $\pm 1^{\circ}\text{C}$ allowed fluctuation. We notice the largest changes in the earliest hours of the test, when the self-correcting system was stabilizing its initial temperature. Even including this initial change, the thermal system provided excellent stability over the 15 hour test. This result, combined with quoted performance capabilities for the heating and insulation, provides confidence that our observatory can maintain the necessary




Figure 5.9 The ambient overnight temperature of the inside of the Hex Tube with all self-sufficient thermal subsystems operating. The probe recording the temperature for this test was located in the rear of the Hex Tube, to measure the stability at the region where the instrument package will be located. During this test, the ambient temperature of the freezer was -15°C , and the internal temperature within the Hex Tube was set to maintain $+30^{\circ}\text{C}$.

temperature stability throughout the Antarctic winter.

5.3.3 Complications for a 0.5m Polar Telescope

Jupiter next rises above an altitude of 10° at the geographic South Pole in November of 2029. It remains observable (above 10° elevation and with the Sun below the horizon) for an estimate of April 15, 2030 to August 30, 2030, for an estimated total of 137 days. This cycle repeats each year during the Polar winter, until May 2034, when Jupiter sets below the horizon once more. Therefore, we have approximately a 4-year cycle, with roughly 100

days per cycle, during which we can observe Jupiter from the South Pole.

Although the reduced spectral confusion due to periodic gaps will be eliminated, this produces only 400 nights of observation total, or a little over one full year. As noted previously, even with a 3.6m telescope, over one *hundred* years will be required to achieve a sensitivity on the order of 0.1 m s^{-1} . Here, we are decreasing our light-collecting area by a factor of *seven* through reduction of the aperture size to only 0.5m, which will drastically increase the duration of time needed to achieve the desired sensitivity.

In particular, we can roughly estimate the additional time increase needed to compensate for decreasing our aperture. We know that we reached a sensitivity S_0 of 4.5 cm s^{-1} on a 24-night time-frame (t_0) with a 21% fill factor (f_0), through a 3.6m telescope diameter (D_0). With this information combined with knowledge of our desired sensitivity of 0.1 cm s^{-1} , our reduced telescope aperture $D_1 = 0.5\text{m}$, and an estimated fill factor of $f_1 = 80\%$ for the South Pole, we can estimate our desired length of time :

$$S_1 = S_0 \frac{D_0 \sqrt{t_0 f_0}}{D_1 \sqrt{t_1 f_1}} \quad (5.9)$$

Rearranging this ratio to solve for the new required time and inputting our modified equipment, we obtain:

$$t_1 = \left(\frac{S_0 D_0}{S_1 D_1} \right)^2 \times \frac{t_0 f_0}{f_1} = \left(\frac{4.5 \text{ cm/s } 3.6 \text{ m}}{0.1 \text{ cm/s } 0.5 \text{ m}} \right)^2 \times \frac{24 \text{ d} \times 0.21}{0.8} = 431,065 \text{ d} \quad (5.10)$$

661,348 days, or 1,810 *years*. There are other ways to improve sensitivity to try and compensate for these losses, such as adding a second channel to increase your spatial coverage

by a factor of approximately 3, reducing this required time to 603 years. Secondly, we could replace the single-lined, narrow bandpass MOFs that we are using with a multi-line molecular MOF, which is currently in the process of being developed. This MOF could have the ability to sense over 100 lines, which would provide a sensitivity gain by approximately a factor of 10, and reducing the required time to reach a sensitivity of 0.1 cm s^{-1} to 60 years—still an exorbitant amount of time from Earth.

Applying post-processing image restoration techniques with ISHWFS data, as discussed prior, can increase our capabilities as well. Wave propagation studies for the polarimetric arm of PMODE reveal that any waves traveling through the atmosphere of Jupiter will grow in size as they move upwards through lower densities, thus displacing larger volumes of particles in the atmosphere, and likely increasing the resultant size of any measured Doppler shift induced by the wave. Therefore, modification of future instrumentation so that it senses higher in the Jovian atmosphere may provide a better chance at detecting any oscillatory modes. Finally, we can look to implementation of Doppler-sensing techniques which are not as photon-limited as MOFs, such as the Mach-Zehnder interferometric techniques implemented by Gaulme et al. (2011), Gonçalves et al. (2019), and the JIVE/JOVIAL team.

Despite these proposed improvements, it is quite possible that reaching this required sensitivity will simply never be attainable from the ground. Our sole saving grace lies in the consideration that the current predicted mode amplitudes for the oscillations are model-based, and we still lack a complete understanding of all processes occurring deep within Jupiter. There is always a chance that current theoretical models are incorrect, and any

oscillatory modes may yet have amplitudes that are both lower than our 4.5 cm s^{-1} noise floor, but larger than predicted at 0.1 cm s^{-1} , and thus, potentially still detectable from the ground. Therefore, we believe that continuing velocimetric monitoring of Jupiter (particularly from Polar sites with extended observational duration) is a worthwhile goal, to continue pushing constraints on the amplitude limit for global modes lower, and to better constrain the atmospheric dynamics, such as zonal winds and atmospheric gravity waves.

5.3.4 A Multi-Site Polar Network for Detecting Jovian Oscillations

In addition to the aforementioned attempts to improve our observational duration and sensitivity, there is one way yet that the ground-based observations can improve beyond the South Pole capabilities. A benefit of LANDIT is that it is small, compact, and transportable—fundamentally, it *must* be, to be delivered and assembled by a small team in a Polar desert at an elevation of 10,000 ft., and the coldest temperatures on the Earth.

This provides one additional benefit: LANDIT can thus be repackaged, deployed to a different location, and subsequently reassembled. As the planets continue their journeys through the sky, Jupiter will eventually set below the South Pole horizon, and will subsequently rise at the North Pole. Theoretically, LANDIT has the capability to remain at the South Pole during the years that Jupiter is in the sky, collecting observations for the entire duration. Once Jupiter sets, the observatory can be relocated to a satisfactory site within the Arctic circle (such as the Thule Air Base in Northern Greenland, or the Polar Environment Atmospheric Research Laboratory (PEARL) base in Northern Canada), where it can begin its observations anew. This will produce a multi-*year* time series, with no nightly periodic

gaps—in particular, after Jupiter sets below the South Pole horizon in May 2034, it rises in the Arctic Circle later that year, in October 2034. Thus, the 4-5 year cycle of collecting Polar observations can begin anew, resulting in ~ 100 consecutive days of data collection with no periodic gaps each year, for theoretically as many years as desired.

The PMODE instrumentation was able to reduce the current state-of-the-art amplitude limits for Jovian seismology to a level of 4.5 cm s^{-1} ; the expansion of this sensitive instrumentation to a multi-site Polar network should provide an even more substantial reduction in amplitude limits. This will expand our knowledge regarding Jovian oscillations and the excitation mechanisms at play within Jupiter from a ground based site, even if reduction of overall noise levels to the predicted amplitudes of 0.1 cm s^{-1} is unachievable from the ground.

5.3.5 The Final Frontier: Detecting Jovian Oscillations from a Dedicated Space-Based Doppler Imager

Finally, as the capabilities of ground-based instrumentation will likely not be able to reach a sensitivity as low as the predicted Jovian oscillations, the next step will be to proceed to the final frontier—launching a dedicated satellite for Jovian seismology into space. This has been proposed in the past (Grasset et al. 2013), but remained unfunded, as it had not yet been evidenced that the oscillations were likely smaller than expected, and perhaps entirely undetectable from a ground-based site.

The most cost-effective option for a space-based Jovian observatory is that of a CubeSat, which could be launched to a Sun-synchronous orbit around the Earth, or to the Lagrange

point L2. This has been discussed by the Jovian interiors and atmospheric dynamics community as a goal to recommend for funding from NASA, with the prime instrument being a Doppler velocimeter—however, as we have evaluated in this chapter, it is likely that single-line MOF Doppler velocimeters will not be able to reach the required sensitivity, as the narrow pass-band results in too substantial of a limitation in photons. Although launching of a CubeSat to space would provide a substantial increase in resolution, a lack of periodic gaps, and higher transmission through fewer optical components, this comes as a result of a drastically downsized aperture, to a maximum of a 10 cm diameter for a typical CubeSat bus. Therefore, the gain in photons from optical transmission will be negated by the decrease in aperture size—in particular, by applying the same order-of-magnitude estimates as in Equation 5.10, decreasing to a 10 cm diameter (while including the increase in signal from expanding the spatial coverage and replacing the single-lined MOF with a multi-lined one), we expect approximately 1,185 years of observation to be necessary from an Earth-orbiting CubeSat to reach a sensitivity of 10 cm s^{-1} .

Finally, there have been historical proposals for a Doppler velocimeter as a primary instrument for a flagship level Jupiter orbiter, which would certainly produce the highest-quality data and the best chance possible to probe the Jovian interior (Grasset et al. 2013). At the time of the original proposal, the Doppler velocimeter was not selected as a science experiment—however, at the time, it was also believed that rock storms were likely the excitation mechanism for the global modes of Jupiter, thereby producing amplitudes up to 50 cm s^{-1} , which could be feasibly measured from ground-based missions—thus, if the goal

was achievable from the ground, it was not necessary to allocate space for the collecting instrument on a flagship-level planetary orbiter mission.

Our work has shown that this is not the case, and instead, for the current model-based predictions, we must reach a drastically lower sensitivity—on the order of 0.1 cm s^{-1} , which is likely so low that the modes will remain undetectable from an Earth-based observatory unless we have the capability to dedicate decades, from a large telescope, solely to this measurement. Similarly, due to the photon restrictions imposed by MOF-based Doppler velocimeters and substantially decreased aperture sizes, the 0.1 cm s^{-1} noise floor required will likely remain unreachable by Earth-orbiting satellites as well. It is very possible that the only attainable way to obtain measurements with a high enough input flux to detect the Jovian oscillations is by placing a dedicated Doppler velocimeter in orbit around Jupiter itself.

This brings us 4.2 AU closer to Jupiter, thus improving our sensitivity by a factor of 4.2^2 as we no longer suffer from the inverse-square falloff of light. Including this factor in Equation 5.10, assuming a full 100% fill factor, a 10 cm entrance aperture to the telescope, a 3x increase in spatial coverage on the Jovian disk, and a 100-line molecular MOF will require only 3.9 years of dedicated observations to reach a sensitivity of 0.1 cm s^{-1} . The Jovian interiors and atmospheric dynamics community has universally agreed that unambiguous detection of the global oscillations is an imperative goal for future leaps in knowledge within this field, and it is becoming clear that a dedicated Jovian orbiter with a Doppler velocimeter may produce the only path forward to meet this goal.

CHAPTER 6

CONCLUSIONS

“It is not so important that you get the right answer: it is most important that you ask the right question.” –Stevenson (2020)

Although the goal we set out to achieve—unambiguous detection of the global modes of Jupiter and clear determination of the interior structure to answer the long-standing questions regarding planetary formation—remains unanswered, the work conducted through this dissertation has significantly progressed the field of Jovian seismology.

We have developed and built two novel sets of instrumentation—PMODE and LAN-DIT—and conducted the longest and most sensitive Doppler measurements of Jupiter to date. This work has countered the previous detection of Jovian global modes with amplitudes of 50 cm s^{-1} , fundamentally changing our understanding of excitation mechanisms at play deep within the Jovian interior. Our work shows no clear excess power greater than 4.5 cm s^{-1} in the prime mode of interest, $Y_\ell^m = (1, 0)$; as such, we suggest that the previously theorized rock storms are not the prime driving source of any oscillatory modes within Jupiter. Instead, it is most likely that the global modes are excited by turbulent convection or water storms, both of which are known to be present within the Jovian atmosphere.

In addition to countering past results regarding the $Y_\ell^m = (1, 0)$ mode, we also provide to the community a thoroughly calibrated and refined report of noise levels for *all* Jovian modes up to and including $\ell = \pm 11$, which we expect to be an invaluable resource as the field grows and progresses. We note that not only were modes not detected in the

$Y_\ell^m = (1, 0)$ mode, but no significant excess power resembling global oscillations has been detected within *any* of these modes, to a maximum sensitivity of $\sim 2 \text{ cm s}^{-1}$. This adds further confidence to our report that global modes are likely on the order of 0.1 cm s^{-1} or below, and future observations from either Polar based sites on Earth (to continue pushing the limit for amplitudes lower) or from a Jupiter-orbiting satellite with a Doppler velocimeter (to truly detect any Jovian global modes) are necessary to further progress the field. We provide our recommendation of utilizing our LANDIT experiment for a Polar-based network for ground-based work to continue expanding our knowledge in the near future, and adapting this small, harsh-environment breadboard for use in a dedicated flagship Jupiter orbiter for Jovian seismology as a long-term goal.

On this journey, we have resolved tension in the field of Jovian atmospheric by producing a new and sensitive measurement of Jupiter’s zonal winds profile. We suggest that the sole previous Doppler measurement does *not*, in fact, contradict the previous well-studied feature-tracking measurements. Instead, we provide clear evidence that our Doppler measurement (sensitive to the high-spatial frequency components), combined with a measurement sensitive to the low-spatial frequency components (such as a filtered version of the feature-tracking measurements, or the previous Doppler measurement), will accurately reproduce and agree with feature tracking models, in both structure and amplitude. This constraint and knowledge on the true velocimetric profile will be invaluable in future modeling of the Jovian interior.

Finally, through this work, we have produced a rich and robust dataset for all of planetary

science. Though this dissertation focused exclusively on our observations of Jupiter, the full collected dataset includes sensitive measurements of the Jovian moons, Mars, Venus, and Uranus. We expect that the capabilities of the PMODE velocimetric dataset for other planetary bodies is sufficient for a multitude of new discoveries to be made, ranging from velocimetric measurements of the zonal wind profile of Uranus, to validation of the presence of a standing gravity wave trapped by the Venusian volcanoes.

For the field of Jupiter interiors and atmospheric dynamics, our Doppler dataset—with centimeter-level velocity sensitivity and exquisite spatial *and* temporal resolution—holds unimaginable possibility for future studies to be conducted from its contents. We expect this dataset can be used to uncover many secrets which currently remain shrouded under the cloudy veils of Jupiter: velocimetric measurements of upwelling signals in the Great Red Spot, new insights on the formation and evolution of storms in the Jovian atmosphere, short-term temporal variability in the Jovian winds, and so on.

As so much potential remains within the observations we have collected through this work, we are currently in the process of freely and openly hosting this PMODE dataset directly alongside previous and future attempts to detect Jovian oscillations, in a complete Jovian Seismology Database hosted by NASA’s Planetary Data Services. We expect our work to be available to the public within the next year—the possibility of what remains to be uncovered in the dataset is limited only by the capability of the minds analyzing it.

“Our curiosity goes with you on your journey. You walk in the footsteps of those who came before you, and your path guides those who will follow later.”

-Outer Wilds

Appendices

A MATLAB Code to Reduce PMODE Data

The MATLAB code in the following pages reads the raw data from the 2020 PMODE observational campaign as an input, and produces reduced, detrotated, and co-registered image stacks for: the MOF images, the continuum images, and the resultant Dopplergrams, as an output for each night.


```

clc
clear

import matlab.io.*

tic

dates = {'8_06','8_07','8_08','8_09','8_10','8_11','8_12','8_13','8_14',
        '8_16','8_17','8_18','8_19','8_20','8_21','8_22','8_23','8_24','8_25',
        '8_26','8_27','8_28','8_30'};

lengthDates = size(dates,2);
n=1;

avg_size = 200;

for n = 1:lengthDates

    tic

    date = dates{n};
    fprintf(strcat('Starting ..',date,'\n'))

    olddatadir = strcat(date,'/K/JUPITER/');
    datadir = strcat(olddatadir,'RENAMED/');
    if ~exist(datadir,'dir')
        mkdir(datadir)
    end

    dataSearch = strcat(olddatadir,'K_JUPITER*.fits');
    files = dir([dataSearch]);

    Temps = [];
    Times = [];
    MJD = [];

    % Set the paths to your data directories from the directory
    % where this code is located -- make sure you put the final slash
    reduceddir = strcat(date,'/K/REDUCED/');
    medBias = fitsread('medBias_K.fits');
    medDark = fitsread('medDark_K.fits');
    medFlat = fitsread('medFlat_K.fits');
    medFlat_bg = fitsread('medFlat_K_background1.fits');
    BPM = fitsread("BPM_K.fits");
    Gradient = fitsread(strcat(reduceddir,date,'-gradient.fits'));

    medLeak = fitsread('medLeak_K.fits');
    medLeak_forscaling = medLeak;
    medLeak_forscaling(medLeak_forscaling < 1000) = 0;
    leak_scaling_idx = find(medLeak_forscaling);
    medLeak(:,265:end) = 0;
    leakmean = nanmean(nanmean(medLeak_forscaling(leak_scaling_idx)));
    flatmean = nanmean(nanmean(medFlat(leak_scaling_idx)));

```

```

medFlat_leakcor = medFlat - (medLeak.*(leakmean/flatmean));

rotRef = fitsread('8_17_rotationRef_K.fits');

if ~exist(reduceddir,'dir')
    mkdir(reduceddir)
end

% Set the strings it should look for to pick the correct filenames
% out of each directory
dataSearch = strcat(datadir,'5*MJD*.fits');

% Make a list with the different filenames you want, using the
% previous information
files = dir([dataSearch]);

% Read in all the actual science light frames into a [size, size, #frame]
% array; also storing the camera temperatures and times for each frame
Data = [];
Data = zeros(513,536,length(files));
Filenames = [];
Data_MJD = [];
Data_Replaced = 0;
Data_Replaced_Idx = [];
Unused_Data = [];
k = 1;
for i = 1:length(files)
    baseFileName = files(i).name;
    Filenames{i} = baseFileName;
    fullFileName = fullfile(datadir,baseFileName);
    test = fitsread(fullFileName);

    if mean(mean(test)) <= 3.0e+03 | mean(mean(test)) > 4e4
        Data(:, :, i) = nan;
        Data_Replaced = Data_Replaced+1;
        Data_Replaced_Idx(k) = i;
        Unused_Data(:, :, k) = test;
        k=k+1;
    else
        Data(:, :, i) = test;
    end

    string=Filenames{i};
    split=strsplit(string,'_');
    Data_MJD(i) = str2num(split{1});
end
fclose('all');

% set constant sizes that'll be used later in the code
sizeData = size(Data,3);
imageSizeY = length(Data(:,1,1));
imageSizeX = length(Data(1,:,1));
vgram_sizeX = 241;
vgram_sizeY = 241;

```

```

vgram_center = round(vgram_sizeX/2);
[columnsInImage, rowsInImage] = meshgrid(1:vgram_sizeX, 1:vgram_sizeY);

% BPM, bias, dark, and flat correction on all frames
parfor i = 1:sizeData
    Data(:, :, i) = Data(:, :, i) .* BPM;
    Data(:, :, i) = fillmissing(Data(:, :, i), 'movmean', 5, 'EndValues', 'nearest');
    Data(:, :, i) = Data(:, :, i) - medBias;
    Data(:, :, i) = Data(:, :, i) - medDark;

    temp = Data(:, :, i);
    conmean = nanmean(nanmean(temp(leak_scaling_idx)));
    scaling_saved(i) = leakmean/conmean;

    Data_Cor(:, :, i) = Data(:, :, i) ./ (medFlat - medFlat_bg);
    Data_LeakCor_Unsmoothed(:, :, i) = (Data(:, :, i) - ...
        (medLeak.*(leakmean/conmean))) ./ (medFlat_leakcor);
end

scaling_saved_smooth = movmedian(scaling_saved, 30, 'omitnan');

parfor i = 1:sizeData
    Data_LeakCor(:, :, i) = (Data(:, :, i) - ...
        (medLeak.*(scaling_saved_smooth(i)))) ./ (medFlat_leakcor);
end

con_x = 374;
con_y = 379;
vel_x = 130;
vel_y = 370;

radius = 120;
diameter = 2*radius;

xmin_con = con_x - radius;
ymin_con = con_y - radius;
xmin_vel = vel_x - radius;
ymin_vel = vel_y - radius;

xshift = -0.6132;
yshift = 1.4628;
saved_conshift_X = zeros(sizeData);
saved_conshift_Y = zeros(sizeData);

vel_saved_leakcor = zeros(241, 241, sizeData);
con_saved_leakcor = zeros(241, 241, sizeData);
vgram_velflatcor_blurr = zeros(241, 241, sizeData);
vgram_velflatcor_blurr_leakcor = zeros(241, 241, sizeData);
edges = zeros(241, 241, sizeData);

intg_intens_vgram = [];
intg_intens_conregion = [];
intg_intens_velregion = [];
diameter_saved = [];

```

```

test_ref = Data_Cor(:, :, 10);
test_ref_idx = 10;
if nanmean(nanmean(test_ref)) == 0 | isnan(nanmean(nanmean(test_ref)))
    test_ref = Data_Cor(:, :, 100);
    test_ref_idx = 100;
    fprintf('yes')
end

conimg_ref = imcrop(test_ref, [xmin_con ymin_con diameter diameter]);
velimg_ref = imcrop(test_ref, [xmin_vel ymin_vel diameter diameter]);
velimg_old_ref = velimg_ref;
velimg_ref = velimg_ref - Gradient;

conimg_ref(isnan(conimg_ref)) = 0;
conimg_ref(isinf(conimg_ref)) = 0;

conReg_ref = imtranslate(conimg_ref, [xshift, yshift]);
vgram_blurr_ref = (velimg_ref ./ conReg_ref);
vgram_blurr_old_ref = (velimg_old_ref ./ conReg_ref);
vgram_blurr_ref(isnan(vgram_blurr_ref)) = 0;
vgram_blurr_old_ref(isnan(vgram_blurr_old_ref)) = 0;

holder_old = vgram_blurr_old_ref;
holder_con = conReg_ref;
test = holder_con;
test(test == 0) = nan;
test = edge(test, 'log', 0.00003);
test_clean = holder_old;
test_clean(test_clean == 0) = nan;
test_clean = edge(test_clean, 'log', 0.03);
test_clean = bwareaopen(test_clean, 100);
test_clean = bwconvhull(test_clean);

stats = regionprops('table', test_clean, 'Centroid', ...
    'Eccentricity', 'EquivDiameter', 'MinorAxisLength', ...
    'MajorAxisLength');
indx = dsearchn(stats.Centroid, [124 124]);
centerX = stats.Centroid(indx, 1);
centerY = stats.Centroid(indx, 2);
radius = (stats.EquivDiameter / 2) - 8;

ellipsePixels = (rowsInImage - centerY).^2 ./ (radius)^2 +
    (columnsInImage - centerX).^2 ./ (radius)^2 <= 1;
test_ref = test .* ellipsePixels;
test_ref_filled = bwconvhull(test_ref);

parfor i = 1:sizeData

    img = Data_Cor(:, :, i);
    img_leakcor = Data_LeakCor(:, :, i);

    if (isnan(nanmean(nanmean(img))))
        vel_saved(:, :, i) = nan;
    end
end

```

```

        con_saved(:, :, i) = nan;
        vgram_velflatcor_blurr(:, :, i) = nan;
        vel_saved_leakcor(:, :, i) = nan;
        con_saved_leakcor(:, :, i) = nan;
        vgram_velflatcor_blurr_leakcor(:, :, i) = nan;
        edges(:, :, i) = nan;
        diameter_saved(i) = nan;
        continue
    end

    % crop out the full field of the velocity image
    % and the continuum image
    conimg = imcrop(img, [xmin_con ymin_con diameter diameter]);
    velimg = imcrop(img, [xmin_vel ymin_vel diameter diameter]);
    conimg_leakcor = imcrop(img_leakcor, [xmin_con ymin_con ...
        diameter diameter]);
    velimg_leakcor = imcrop(img_leakcor, [xmin_vel ymin_vel ...
        diameter diameter]);

    velimg_old = velimg;

    conimg(isnan(conimg)) = 0;
    conimg(isinf(conimg)) = 0;
    conimg_leakcor(isnan(conimg_leakcor)) = 0;
    conimg_leakcor(isinf(conimg_leakcor)) = 0;

    conReg = imtranslate(conimg, [xshift, yshift], 'bilinear', ...
        'OutputView', 'same');
    conReg_leakcor = imtranslate(conimg_leakcor, [xshift, yshift], ...
        'bilinear', 'OutputView', 'same');

    vgram_blurr = (velimg ./ conReg);
    vgram_blurr_old = (velimg_old ./ conReg);
    vgram_blurr_leakcor = (velimg_leakcor ./ conReg_leakcor);
    vgram_blurr(isnan(vgram_blurr)) = 0;
    vgram_blurr_old(isnan(vgram_blurr_old)) = 0;
    vgram_blurr_leakcor(isnan(vgram_blurr_leakcor)) = 0;

    holder_old = vgram_blurr_old;
    holder_con = conReg;
    test = holder_con;
    test(test == 0) = nan;
    test = edge(test, 'log', 0.00003);
    test_clean = holder_old;
    test_clean(test_clean == 0) = nan;
    test_clean = edge(test_clean, 'log', 0.03);
    test_clean = bwareaopen(test_clean, 100);
    test_clean = bwconvhull(test_clean);

    stats = regionprops('table', test_clean, 'Centroid', ...
        'Eccentricity', 'EquivDiameter', 'MinorAxisLength', ...
        'MajorAxisLength');
    indx = dsearchn(stats.Centroid, [124 124]);
    centerX = stats.Centroid(indx, 1);

```

```

centerY = stats.Centroid(indx,2);
radius = (stats.EquivDiameter/2) - 8;
diameter_saved(i) = stats.MajorAxisLength;

ellipsePixels = (rowsInImage - centerY).^2 ./ (radius)^2
+ (columnsInImage - centerX).^2 ./ (radius)^2 <= 1;
test = test.*ellipsePixels;
test_safe = test;
test = bwareaopen(test,50);
test_edge = test;
test = bwconvhull(test);

[idx_row, idx_col] = find(test);
xmin_indx = min(idx_col);
xmax_indx = max(idx_col);
ymin_indx = min(idx_row);
ymax_indx = max(idx_row);

if (length(xmin_indx) == 0) | (length(xmax_indx) == 0)
    | (length(ymin_indx) == 0) | (length(ymax_indx) == 0)
    test = test_safe;
    test = bwareaopen(test,100);
    test_edge = test;
    test = bwconvhull(test_edge);
end

if length(find(test)) < 15000
    test = test_safe;
    test = bwareaopen(test,100);
    test_edge = test;
    test = bwconvhull(test_edge);
end

bckgrnd_indx = find(~test);

vgram_blurr(bckgrnd_indx) = 0;
vgram_blurr_leakcor(bckgrnd_indx) = 0;
velimg_leakcor(bckgrnd_indx) = 0;
conReg_leakcor(bckgrnd_indx) = 0;

stats = regionprops('table', test, 'Centroid', 'Eccentricity', ...
    'EquivDiameter', 'MinorAxisLength', 'MajorAxisLength');
indx = dsearchn(stats.Centroid,[124 124]);
centerX = stats.Centroid(indx,1);
centerY = stats.Centroid(indx,2);

velimg_leakcor = imtranslate(velimg_leakcor,[120-centerX, 120-centerY]);
conReg_leakcor = imtranslate(conReg_leakcor,[120-centerX, 120-centerY]);

saved_conshift_X(i) = 120-centerX;
saved_conshift_Y(i) = 120-centerY;

vgram_blurr(isnan(vgram_blurr)) = 0;
vgram_blurr(isinf(vgram_blurr)) = 0;

```

```

vgram_blurr_leakcor(isnan(vgram_blurr_leakcor)) = 0;
vgram_blurr_leakcor(isinf(vgram_blurr_leakcor)) = 0;

vgram_blurr = imtranslate(vgram_blurr,[120-centerX, 120-centerY], ...
    'bilinear','OutputView','same');
vgram_blurr_leakcor = imtranslate(vgram_blurr_leakcor, ...
    [120-centerX, 120-centerY],'bilinear','OutputView','same');

test_edge(isnan(test_edge)) = 0;
test_edge(isinf(test_edge)) = 0;
test_edge = real(test_edge);

test_edge = imtranslate(test_edge,[120-centerX, 120-centerY], ...
    'bilinear','OutputView','same');

edges(:,:,i) = test_edge;
vel_saved_leakcor(:,:,i) = velimg_leakcor;
con_saved_leakcor(:,:,i) = conReg_leakcor;

vgram_velflatcor_blurr(:,:,i) = vgram_blurr;
vgram_velflatcor_blurr_leakcor(:,:,i) = vgram_blurr_leakcor;

end

test_ref = vgram_velflatcor_blurr(:,:,10);
test_ref_idx = 10;
if nanmean(nanmean(test_ref)) == 0 | isnan(nanmean(nanmean(test_ref)))
    test_ref = vgram_velflatcor_blurr(:,:,100);
    test_ref_idx = 100;
    fprintf('yes')
end

% Derotate the Dopplergrams

EPHEMERIS = readtable('horizons_resultsFULLRUN.txt') ;
EPHEMERIS=removevars(EPHEMERIS,"Var2");
EPHEMERIS=removevars(EPHEMERIS,"Var3");
EPHEMERIS = EPHEMERIS{:,:};
%Grab site info
lat=20.7083119;

%%Timing info
EPHEMERIS_TIME=EPHEMERIS(:,1)-2400000.5;
az=EPHEMERIS(:,4)-180;
alt=EPHEMERIS(:,5);
dec=EPHEMERIS(:,3);
ha=EPHEMERIS(:,8);
ParaAngle=[];
ha2=[];

%Calculate the PA (15 on ha is for hour angle to degrees)
for i = 1:length(az)
    ParaAngle(i)=atan2d(sind(ha(i)*15),(tand(lat).*cosd(dec(i)))- ...
        (sind(dec(i)).*cosd(ha(i)*15)));

```

```

end
i=1;

%Grab the position angle for Jupiter
PosAngle=EPHEMERIS(:,6)';

%Rotation calc
RotateNorth=[];
ParaAngle=ParaAngle';
PosAngle=PosAngle';
RotateNorth=abs(ParaAngle-az+alt-PosAngle+29.5);

vgramRot = zeros(241,241,sizeData);
vgramRot_leakcor = zeros(241,241,sizeData);
conRot_leakcor = zeros(241,241,sizeData);
velRot_leakcor = zeros(241,241,sizeData);
edgesRot = zeros(241,241,sizeData);
IndexRotation= zeros(sizeData);
i=1;
%%Find closest time to image that is being read and rotate
parfor i = 1:sizeData
    [~,IndexRotation2] = (min(abs(EPHEMERIS_TIME - Data_MJD(i))));
    IndexRotation(i)=IndexRotation2;

    vgramRot_temp = imresize(vgram_velflatcor_blurr(:, :, i), 4);
    vgramRot_temp = imrotate(vgramRot_temp, RotateNorth(IndexRotation(i)), ...
        'nearest', 'crop');
    vgramRot(:, :, i) = imresize(vgramRot_temp, 0.25);

    vgramRot_leakcor_temp = imresize(vgram_velflatcor_blurr_leakcor(:, :, i), 4);
    vgramRot_leakcor_temp = imrotate(vgramRot_leakcor_temp, ...
        RotateNorth(IndexRotation(i)), 'nearest', 'crop');
    vgramRot_leakcor(:, :, i) = imresize(vgramRot_leakcor_temp, 0.25);

    conRot_leakcor_temp = imresize(con_saved_leakcor(:, :, i), 4);
    conRot_leakcor_temp = imrotate(conRot_leakcor_temp, ...
        RotateNorth(IndexRotation(i)), 'nearest', 'crop');
    conRot_leakcor(:, :, i) = imresize(conRot_leakcor_temp, 0.25);

    velRot_leakcor_temp = imresize(vel_saved_leakcor(:, :, i), 4);
    velRot_leakcor_temp = imrotate(velRot_leakcor_temp, ...
        RotateNorth(IndexRotation(i)), 'nearest', 'crop');
    velRot_leakcor(:, :, i) = imresize(velRot_leakcor_temp, 0.25);

    edgesRot_temp = imresize(edges(:, :, i), 4);
    edgesRot_temp = imrotate(edgesRot_temp, RotateNorth(IndexRotation(i)), ...
        'nearest', 'crop');
    edgesRot(:, :, i) = imresize(edgesRot_temp, 0.25);

end

test_ref_rot = vgramRot_leakcor(:, :, test_ref_idx);
test_ref_edge_filled = imbinarize(test_ref_rot);

```



```

stats = regionprops('table', test_ref_edge_filled, 'Centroid', ...
    'Eccentricity', 'EquivDiameter', 'MinorAxisLength', 'MajorAxisLength');
indx = dsearchn(stats.Centroid,[120 120]);
centerX = stats.Centroid(indx,1);
centerY = stats.Centroid(indx,2);

test_ref_edge_filled = imtranslate(test_ref_edge_filled, ...
    [120-centerX, 120-centerY], 'nearest', 'OutputView', 'same');
test_ref_edge_filled = double(test_ref_edge_filled);

vgramRot_shifted = zeros(241,241,sizeData);
vgramRot_shifted_leakcor = zeros(241,241,sizeData);
conRot_shifted_leakcor = zeros(241,241,sizeData);
velRot_shifted_leakcor = zeros(241,241,sizeData);
edges = zeros(241,241,sizeData);

parfor i = 1:sizeData

    holder = vgramRot(:,:,i);
    holder_leakcor = vgramRot_leakcor(:,:,i);
    holder_vel_leakcor = velRot_leakcor(:,:,i);
    holder_con_leakcor = conRot_leakcor(:,:,i);
    test_edge = edgesRot(:,:,i);

    if nansum(nansum(holder)) == 0
        vgramRot_shifted(:,:,i) = nan;
        vgramRot_shifted_leakcor(:,:,i) = nan;
        conRot_shifted_leakcor(:,:,i) = nan;
        velRot_shifted_leakcor(:,:,i) = nan;
        edges(:,:,i) = nan;
        continue
    end

    test = double(imbinarize(holder_leakcor));
    movingReg = registerImages(test,test_ref_edge_filled,holder);
    vgram = movingReg.RegisteredImage;

    vgram_leakcor = imwarp(holder_leakcor, ...
        movingReg.Transformation, 'OutputView', imref2d(size(vgram)), ...
        'interp', 'nearest');
    holder_vel_leakcor = imwarp(holder_vel_leakcor,movingReg.Transformation, ...
        'OutputView', imref2d(size(vgram)), 'interp', 'nearest');
    holder_con_leakcor = imwarp(holder_con_leakcor,movingReg.Transformation, ...
        'OutputView', imref2d(size(vgram)), 'interp', 'nearest');
    test_edge = imwarp(test_edge,movingReg.Transformation, 'OutputView', ...
        imref2d(size(vgram)), 'interp', 'nearest');

    vgram(vgram == 0) = nan;
    vgram_leakcor(vgram_leakcor == 0) = nan;
    holder_vel_leakcor(holder_vel_leakcor == 0) = nan;
    holder_con_leakcor(holder_con_leakcor == 0) = nan;

    vgramRot_shifted(:,:,i) = vgram;
    vgramRot_shifted_leakcor(:,:,i) = vgram_leakcor;

```

```

        conRot_shifted_leakcor(:,:,i) = holder_con_leakcor;
        velRot_shifted_leakcor(:,:,i) = holder_vel_leakcor;
        edges(:,:,i) = test_edge;

    end

    fprintf("Registration done, now writing to files...")

    fitwrite(vgramRot_shifted_leakcor, strcat(reduceddir, ...
        date, '-vgram_Rot_070522_leakcor.fits'))
    fitwrite(velRot_shifted_leakcor, strcat(reduceddir, ...
        date, '-vel_Rot_070522.fits'))
    fitwrite(conRot_shifted_leakcor, strcat(reduceddir, ...
        date, '-con_Rot_070522.fits'))
    save(strcat(reduceddir, date, '-leakscaling.mat'), 'scaling_saved')

    toc

end

fprintf("Data reduction finished!")
toc

```

B MATLAB Code to Generate PMODE Residuals

The MATLAB code in the following pages reads the reduced data from the code in Appendix A as an input, calculates and subtracts the average Dopplergram, restricts the Dopplergram to only the velocity sensitive regime, applies the “data number to velocimetric” MOF unit conversion, and produces single stack of residual Dopplergrams for each night as the output.

```

clc
clear

import matlab.io.*
format long

dates = {'8_06','8_07','8_08','8_09','8_10','8_11','8_12','8_13','8_14',
        '8_15','8_16','8_17','8_18','8_19','8_20','8_21','8_22','8_23','8_24',
        '8_25','8_26','8_27','8_28','8_30'};

lengthDates = size(dates,2);

n=1;
avg_size = 200;

RadialCorrection_array = fitsread('RadialCorrectionStack.fits');
MOF_Units = fitsread('MOF_Units_FullImg.fits');
cut_indxs = [524,470,742,731,741,747,760,793,491,761,488,768,765,50,550,
             673,673,500,228,638,736,617,610,591];

tic

for n = 1:lengthDates

    tic

    date = dates{n};
    fprintf(strcat('Starting ..',date,'\n'))

    reduceddir = strcat(date,'/K/REDUCED/');
    fullFileName = strcat(reduceddir,date,'-vgram_Rot_070522_leakcor.fits');
    vgrams = fitsread(fullFileName);

    if strcmp(date,'8_28')
        vgrams(:, :, 123) = nan;
        vgrams(:, :, 339) = nan;
    end

    RadialCorrection = RadialCorrection_array(:, :, n);
    MOF_Unit = MOF_Units(:, :, n);

    MOF_Unit_Con = MOF_Unit;
    [mof_rows, mof_columns] = size(MOF_Unit_Con);
    for row = 1 : mof_rows
        % Extract just this column only.
        thisRow = MOF_Unit_Con(row, :);
        % Shift it down "col" rows.
        shiftedRow = circshift(thisRow, 75);
        % Put it back in
        MOF_Unit_Con(row, :) = shiftedRow;
    end

    oldresiduals_name = strcat('Residuals/0221_ConstAvg/', ...

```

```

    date, '-prelim_residuals_070522.mat');
oldresiduals = load([oldresiduals_name]);
MJD = oldresiduals.MJD_residuals(:,1);
old_residuals = oldresiduals.MJD_residuals(:,2);

test_ref = vgrams(:, :, 10);
test_ref_idx = 10;
if nanmean(nanmean(test_ref)) == 0 | isnan(nanmean(nanmean(test_ref)))
    test_ref = vgrams(:, :, 100);
    test_ref_idx = 100;
    fprintf('yes')
end

test_ref(isnan(test_ref)) = 0;
test_ref(isinf(test_ref)) = 0;

sz = size(vgrams(:, :, 1));
vgram_sizeX = size(test_ref, 1);
vgram_sizeY = size(test_ref, 2);
vgram_center = round(vgram_sizeX/2);
sizevgrams = size(vgrams, 3);

[columnsInImage rowsInImage] = meshgrid(1:vgram_sizeX, 1:vgram_sizeY);

%-----
cut_idx = cut_idx(n);
vgrams(:, :, cut_idx:end) = nan;

full_avg_newvgram = nanmedian(vgrams, 3);
fitswrite(full_avg_newvgram, strcat('ZonalWindAves/', ...
    date, '-avg_fulldisk.fits'));
avgcut = nanmedian(full_avg_newvgram(111:130, 0+25:vgram_sizeX-50));
avgcut_full = nanmedian(full_avg_newvgram(:, 1:vgram_sizeX));
avgcut_full(:, 1:22) = nan;
avgcut_full(:, vgram_sizeX-25:end) = nan;
avgcut_full = avgcut_full.*-1;

avgcut(avgcut >= 3) = nan;
[min_val min_idx] = min(avgcut);
cutoff = min(avgcut) + 0.80*(nanmedian(nanmedian(full_avg_newvgram)) ...
    - min(avgcut));

vel = full_avg_newvgram;
clean_avg_newvgram = full_avg_newvgram;
con = full_avg_newvgram;
clean_avg_newvgram_con = full_avg_newvgram;

test_radialmask = full_avg_newvgram;
test_radialmask(isnan(test_radialmask)) == 0;
test_radialmask = double(imbinarize(test_radialmask));
[idx_row idx_col] = find(test_radialmask);
xmin_idx = min(idx_col);
xmax_idx = max(idx_col);
ymin_idx = min(idx_row);

```

```

ymax_indx = max(idx_row);
centerX = xmin_indx + round(((xmax_indx - xmin_indx)/2));
centerY = ymin_indx + round(((ymax_indx - ymin_indx)/2));

% setting how much of the external diameter we want to crop off
xdiameter = 0.90*(xmax_indx - xmin_indx);
ydiameter = 0.90*(ymax_indx - ymin_indx);
xradius = ceil(xdiameter/2);
yradius = ceil(ydiameter/2);
ellipsePixels = (rowsInImage - centerY).^2 ./ (yradius)^2
+ (columnsInImage - centerX).^2 ./ (xradius)^2 <= 1;

vel(vel > cutoff) = 0;
idx = isnan(vel);
vel(idx) = 0;
vel(:,1:25+(min_indx - 27)) = 0;
vel(:,25+(min_indx + 27):end) = 0;

vel = bwareaopen(vel,1000);
[vel_idx_row vel_idx_col] = find(vel);
vel(:,25+min_indx-3:25+min_indx+3) = 0;
vel = vel.*ellipsePixels;

vel_idx = find(vel);
background_idx = find(~vel);
clean_avg_newvgram(background_idx) = NaN;

[vel_idx_row vel_idx_col] = find(vel);

con_idx_col = vel_idx_col + 75;
con_idx_row = vel_idx_row;
con_idx = sub2ind(sz,con_idx_row,con_idx_col);
blank = zeros(sz);
blank(con_idx) = 1;
background_idx_con = find(~blank);
clean_avg_newvgram_con(background_idx_con) = NaN;

fitswrite(clean_avg_newvgram, strcat(reduceddir,date, ...
    '-avg_vgram_radialcrop_leakcor.fits'));
fitswrite(clean_avg_newvgram_con, strcat(reduceddir,date, ...
    '-avg_vgram_radialcrop_leakcor_con.fits'));

residuals_nospharm = zeros([1 sizevgrams]);
residuals_nospharm_con = zeros([1 sizevgrams]);
residual_imgs_nospharm = zeros(vgram_sizeX,vgram_sizeY,sizevgrams);
residual_imgs_nospharm_con = zeros(vgram_sizeX,vgram_sizeY,sizevgrams);

true_min_indx = min_indx + 25;
true_min_indx_con = true_min_indx + 75;

parfor i = 1:sizevgrams

    singlemasked = vgrams(:, :, i);

```

```

if (isnan(nanmean(nanmean(singlemasked))))

    residual_imgs_nospharm(:, :, i) = nan;
    residual_imgs_nospharm_con(:, :, i) = nan;
    residuals_nospharm(i) = nan;
    residuals_nospharm_con(i) = nan;
    continue

end

singlelevel = singlemasked;
singlelevel(background_idx) = nan;
singlelevel = singlelevel - clean_avg_newvgram;

singlelevel = singlelevel./MOF_Unit;
singlelevel_std = nanstd(singlelevel(:));
singlelevel_mean = nanmedian(nanmedian(singlelevel));
singlelevel(singlelevel > (singlelevel_mean+(2*singlelevel_std))) = nan;
singlelevel(singlelevel < (singlelevel_mean-(2*singlelevel_std))) = nan;
singlelevel(:, 1:true_min_idx) = singlelevel(:, 1:true_min_idx).*(-1);
residual_imgs_nospharm(:, :, i) = singlelevel;
residuals_nospharm(i) = nanmean(nanmean(singlelevel));

%%% continuum below %%%

singlecon = singlemasked;
singlecon(background_idx_con) = nan;
singlecon = singlecon - clean_avg_newvgram_con;
singlecon = singlecon.*RadialCorrection;
singlecon = singlecon./MOF_Unit_Con;
singlecon_std = nanstd(singlecon(:));
singlecon_mean = nanmean(nanmean(singlecon));
singlecon(singlecon > (singlecon_mean+(2*singlecon_std))) = nan;
singlecon(singlecon < (singlecon_mean-(2*singlecon_std))) = nan;
singlecon(:, 1:true_min_idx_con) = singlecon(:, 1:true_min_idx_con).*(-1);
residual_imgs_nospharm_con(:, :, i) = singlecon;
residuals_nospharm_con(i) = nanmean(nanmean(singlecon));

end

residuals_nospharm_smooth = residuals_nospharm -
    smoothdata(residuals_nospharm, 'rloess', 200);
residuals_nospharm_smooth_con = residuals_nospharm_con -
    smoothdata(residuals_nospharm_con, 'rloess', 200);

MJD_residuals = [MJD residuals_nospharm'];
MJD_residuals_con = [MJD residuals_nospharm_con'];

save(strcat('RESIDUALS_K_CLEAN/NO_SPHARM/', date, ...
    '-residuals_nospharm_scaled_NORADIAL.mat'), 'MJD_residuals');
save(strcat('RESIDUALS_K_CLEAN/NO_SPHARM/', date, ...
    '-residuals_nospharm_scaled_con_NORADIAL.mat'), 'MJD_residuals_con');

fitwrite(residual_imgs_nospharm, ...

```

```
        strcat('RESIDUALS_K_CLEAN/NO_SPHARM/',date, ...
        '-residuals_nospharm_scaled_NORADIAL.fits'));
    fitswrite(residual_imgs_nospharm_con, ...
    strcat('RESIDUALS_K_CLEAN/NO_SPHARM/',date, ...
    '-residuals_nospharm_scaled_con_NORADIAL.fits'));

    toc

end

toc
```


C MATLAB Code to Multiply Residuals by Spherical Harmonics

The MATLAB code in the following pages reads the residual images from the code in Appendix B as an input, then multiplies each image for each night by all spherical harmonics up to $\ell = 11$ and all corresponding m values, producing one stack, per spherical harmonic, per night as the output.

```

clc
clear

import matlab.io.*
format long

dates = {'8_06','8_07','8_08','8_09','8_10','8_11','8_12','8_13','8_14',
        '8_15','8_16','8_17','8_18','8_19','8_20','8_21','8_22','8_23','8_24',
        '8_25','8_26','8_27','8_28','8_30'};

lengthDates = size(dates,2);
fileindx = 1;

n=1;

tic

lrange = 0:11;

for n = 1:lengthDates

    tic

    date = dates{n};
    fprintf(strcat('Starting ..',date,'\n'))

    fullFileName = strcat('RESIDUALS_K_CLEAN/NO_SPHHARM/',date, ...
        '-residuals_nospharm_scaled_NORADIAL.fits');
    residual_imgs = fitsread(fullFileName);

    fullFileName_con = strcat('RESIDUALS_K_CLEAN/NO_SPHHARM/',date, ...
        '-residuals_nospharm_scaled_con_NORADIAL.fits');
    residual_imgs_con = fitsread(fullFileName_con);

    oldresiduals_name = strcat('RESIDUALS_K_CLEAN/NO_SPHHARM/',date, ...
        '-residuals_nospharm_scaled.mat');
    oldresiduals = load([oldresiduals_name]);
    MJD = oldresiduals.MJD_residuals(:,1);

    sizeresiduals = size(residual_imgs,3);

    residuals_sphharm = zeros([1 sizeresiduals]);
    residuals_sphharm_con = zeros([1 sizeresiduals]);
    residual_imgs_sphharm = zeros(241,241,sizeresiduals);
    residual_imgs_sphharm_con = zeros(241,241,sizeresiduals);

    for j = 1:length(lrange)

        l = lrange(j);
        mrange = -1:1;

        for k = 1:length(mrange)

```

```

m=mrangle(k);

fprintf(strcat('Starting sphharm l,m= ',num2str(l),',' ...
,num2str(m),'\n'))

sphharm_array_name = strcat('SphHarmArrays/SphHarm_', ...
    num2str(l), '_', num2str(m), '.fits');
sphharm_array = fitsread(sphharm_array_name);
sphharm = sphharm_array(:,:,n);

parfor i = 1:sizeresiduals

    residual_img = residual_imgs(:,:,i);
    residual_img_con = residual_imgs_con(:,:,i);

    if (isnan(nanmean(nanmean(residual_img))))

        residual_images_sphharm(:,:,i) = 0;
        residual_images_sphharm_con(:,:,i) = 0;

        residuals_sphharm(i) = nan;
        residuals_sphharm_con(i) = nan;

        continue

    end

    residual_img_sphharm = residual_img.*sphharm;
    residual_imgs_sphharm(:,:,i) = residual_img_sphharm;
    residuals_sphharm(i) = nanmean(nanmean(residual_img_sphharm));

    %% CONTINUUM BELOW %%

    residual_img_sphharm_con = residual_img_con.*sphharm;
    residual_imgs_sphharm_con(:,:,i) = residual_img_sphharm_con;
    residuals_sphharm_con(i) = nanmean(nanmean( ...
        residual_img_sphharm_con));

end

MJD_residuals = [MJD residuals_sphharm'];
MJD_residuals_con = [MJD residuals_sphharm_con'];

datadir = strcat('RESIDUALS_K_CLEAN/', num2str(l), '_', num2str(m), '/');
if ~exist(datadir, 'dir')
    mkdir(datadir)
end

mult_sphharm_residuals_filename = strcat('RESIDUALS_K_CLEAN/', ...
    num2str(l), '_', num2str(m), '/', date, ...
    '-residuals_sphharm_scaled_NORADIAL.mat');
mult_sphharm_residuals_filename_con = strcat('RESIDUALS_K_CLEAN/', ...
    num2str(l), '_', num2str(m), '/', date, ...
    '-residuals_sphharm_scaled_con_NORADIAL.mat');

```

```

save(mult_sphharm_residuals_filename,'MJD_residuals');
save(mult_sphharm_residuals_filename_con,'MJD_residuals_con');

mult_sphharm_images_filename = strcat('RESIDUALS_K_CLEAN/', ...
    num2str(l), '_', num2str(m), '/', date, ...
    '-residuals_sphharm_scaled_NORADIAL.fits');
mult_sphharm_images_filename_con = strcat('RESIDUALS_K_CLEAN/', ...
    num2str(l), '_', num2str(m), '/', date, ...
    '-residuals_sphharm_scaled_con_NORADIAL.fits');

fitswrite(residual_imgs_sphharm,mult_sphharm_images_filename);
fitswrite(residual_imgs_sphharm_con,mult_sphharm_images_filename_con);

toc

end

end

end

toc

```

D MATLAB Code to Calculate Final Amplitude Spectra

The MATLAB code in the following pages reads the residual images multiplied by the respective, desired spherical harmonic from the code in Appendix C as an input, then combines each night into a full time series, and subsequently calculates the Lomb-Scargle periodogram to obtain the final amplitude spectra as an output.

```
clc
clear
```

```
% Set desired date range
```

```
dates = {'8_06','8_07','8_08','8_09','8_10','8_11','8_12','8_13','8_14',...
         '8_15','8_16','8_17','8_18','8_19','8_20','8_21',...
         '8_22','8_23','8_24','8_25','8_26','8_27','8_28','8_30'};
```

```
lengthdates=length(dates);
```

```
lrange = 6:10;
```

```
lrange = 1;
```

```
%% Read in data from each night
```

```
for j = 1:length(lrange)
```

```
    l = lrange(j);
```

```
    mrange = -1:1;
```

```
    for k = 1:length(mrange)
```

```
        m=mrange(k);
```

```
        fprintf(strcat('Starting spharm l,m= ',num2str(l),',', ...
                        num2str(m),'\n'))
```

```
        base_spharmdir_name = strcat('RESIDUALS_K_CLEAN/', ...
                                      num2str(l),'_',num2str(m),'/');
```

```
        for n = 1:lengthdates
```

```
            date = dates{n};
```

```
            fullFileName = strcat(base_spharmdir_name,date, ...
                                  '-residuals_spharm_scaled_NORADIAL.mat');
```

```
            Data{n} = load(fullFileName);
```

```
        end
```

```
%% Grab raw data and time vector
```

```
for i=1:lengthdates
```

```
    holder = [Data{i}.MJD_residuals(1:end-1,2)];
```

```
    dataRAW{i}=[holder];
```

```
    dataFSmooth = smoothdata(holder,'rloess',200);
```

```
    holder_smooth=(holder-dataFSmooth);
```

```
    temp = find(~isnan(holder_smooth));
```

```
    if length(temp) == 0
```

```
        last_non_nan_indx = length(holder_smooth);
```

```
        holder_smooth_temp = holder_smooth(1:last_non_nan_indx);
```

```
        tukeyholder = holder_smooth_temp.*
```

```
        tukeywin(length(holder_smooth_temp),.1);
```

```

        tukeyholder(last_non_nan_idx+1:length(holder_smooth)) = nan;
        dataFULL{i}=[tukeyholder];
        Time{i}=[Data{i}.MJD_residuals(1:end-1,1)];
        continue
    end

    last_non_nan_idx = temp(end);
    holder_smooth_temp = holder_smooth(1:last_non_nan_idx);
    tukeyholder = holder_smooth_temp.*
        tukeywin(length(holder_smooth_temp),.1);
    tukeyholder(last_non_nan_idx+1:length(holder_smooth)) = nan;
    dataFULL{i}=[tukeyholder];
    Time{i}=[Data{i}.MJD_residuals(1:end-1,1)];
end

j2=1;
%Combine time vectors
for i=1:lengthdates
    lengthin=length(Time{i});
    for j=1:lengthin
        data_FULL(j2)=[dataFULL{i}(j)];
        TimeSquash(j2)=[Time{i}(j)];
        j2=j2+1;
    end
end

data_FULL=[data_FULL]';

%% Put the correct spacing in the TIME array
time=TimeSquash;
datetime.setDefaultFormats('default','HH:mm:ss yyyy-MM-dd');
dta = datetime(time, 'ConvertFrom','modifiedjuliandate', ...
    'Format','MMM/dd/yyyy HH:mm:ss');

dateTimeShifted_safe = dta;

dateTimeShifted = dateshift(dta,'start','minute');
minuteValues = second(dta);
dateTimeShifted(minuteValues >= 15 & minuteValues < 45)
    = dateTimeShifted(minuteValues >= 15 & minuteValues < 45)
    + seconds(30);
dateTimeShifted(minuteValues >= 45) =
    dateTimeShifted(minuteValues >= 45) + minutes(1);
data=TimeSquash;

%Retime
dataRAW=TimeSquash;
timeRAW=time;
tt=timetable(dateTimeShifted,data');
tv = (dateTimeShifted(1):minutes(0.5):dateTimeShifted(1)+minutes(1)).';
newtt = retime(tt, 'secondly','fillwithmissing');
tv = (newtt.Time(1):minutes(0.5):dateTimeShifted(end)+minutes(1)).';
newtt = retime(tt, tv, 'fillwithmissing');
TimeSquash_Gaps=newtt.Var1;

```

```

%% Put the correct spacing in the VELOCITY array

time=TimeSquash;
datetime.setDefaultFormats('default','HH:mm:ss yyyy-MM-dd');
dta = datetime(time, 'ConvertFrom','modifiedjuliandate', ...
    'Format','MMM/dd/yyyy HH:mm:ss');

dta = dateTimeShifted_safe;
dateTimeShifted = dateshift(dta,'start','minute');
minuteValues = second(dta);
dateTimeShifted(minuteValues >= 15 & minuteValues < 45) =
    dateTimeShifted(minuteValues >= 15 & minuteValues < 45)
    + seconds(30);
dateTimeShifted(minuteValues >= 45) =
    dateTimeShifted(minuteValues >= 45) + minutes(1);
data=data_FULL;
%Retime
dataRAW=data_FULL;
timeRAW=time;
tt=timetable(dateTimeShifted',data);
tv = (dateTimeShifted(1):minutes(0.5):dateTimeShifted(1) ...
    +minutes(1)).';
newtt = retime(tt, 'secondly','fillwithmissing');
tv = (newtt.Time(1):minutes(0.5):dateTimeShifted(end)+minutes(1)).';
newtt = retime(tt, tv, 'fillwithmissing');
data_FULL=newtt.data;

time=newtt.Time;
timespace=time;

%% OUTLIER REJECTION
thresh=4;
outlierlength=15200;

%dataFULL_Raw=data_FULL;
smoothedDATA = data_FULL;

[SmoothCleanedData,outlierIndices,thresholdLow,thresholdHigh] = ...
    filloutliers(smoothedDATA,nan,'movmedian',outlierlength, ...
    'ThresholdFactor',thresh);
data_FULL=SmoothCleanedData;

```

```

v = data_FULL;
t = TimeSquash_Gaps;

t = t.*24.*60.*60;
t = t - t(1);

fact = 3;

```



```

a = 1;
Fc = .00500; % hertz
x = a*sin(2*pi*Fc*t);

% uncomment this if you want to add in a calibration peak
% v = v + x;

[amp,freq] = comp_spectrum(v,t,[0,0.006],'scarg','amp',fact);
[amp_x,freq_x] = comp_spectrum(x,t,[0,0.006],'scarg','amp',fact);

```

```

data_ts = [TimeSquash_Gaps,data_FULL];
timeseries_name = strcat('Results/TimeSeries/timeseries_',num2str(1),'_',num2str(m),'.c');
dlmwrite(timeseries_name, data_ts, 'delimiter', ',', 'precision', 15);

data_ps = [freq.*1000000,amp];
powerspec_name = strcat('Results/PowerSpectra/powerspec_',num2str(1),'_',num2str(m),'.c');
dlmwrite(powerspec_name, data_ps, 'delimiter', ',', 'precision', 15);

```

```

% difference imaging test

diffv = diff(v);
diffx = diff(x);
diff_t = t(2:end);

[diff_amp,diff_freq] = comp_spectrum(diffv,diff_t,[0,0.006], ...
    'scarg','amp',fact);
[diff_amp_x,diff_freq_x] = comp_spectrum(diffx,diff_t,[0,0.006], ...
    'scarg','amp',fact);

diffimg_powerspec_filename = strcat(['Results/' ...
    'DiffImg_PowerSpec_Figures/diffimg_powerspec_'],num2str(1), ...
    ' ',num2str(m),'.png');
figure
plot(diff_freq.*1000000,diff_amp)
xlim([0 6000])
xlabel('Frequency (uHz)')
ylabel('Amplitude (m/s)')
title('Difference Imaging Amplitude Spectrum')
saveas(gcf,diffimg_powerspec_filename);

```

```

timeseries_filename = strcat('Results/TimeSeries_Figures/timeseries_', ...
    num2str(1),'_',num2str(m),'.png');
powerspec_filename = strcat('Results/PowerSpectra_Figures/powerspec_', ...
    num2str(1),'_',num2str(m),'.png');

figure
plot(TimeSquash_Gaps-TimeSquash_Gaps(1),data_FULL+.6)
xlabel('Days')
ylabel('Velocity (m/s)')
title(strcat('l,m= ',num2str(1),' ',num2str(m),[' Time Series ' ...

```

```

        '(smoothed & outlier rejected')]))
ylim([-50 50])
saveas(gcf,timeseries_figname);

figure
plot(freq.*1000000,amp)
xlim([0 6000])
xlabel('Frequency ( $\mu\text{Hz}$ )','Interpreter','latex')
ylabel('Amplitude (m s $^{-1}$ )','Interpreter','latex')
title(strcat('l,m= ',num2str(l),', ',num2str(m),' Power Spectrum'))
saveas(gcf,powerspec_figname);

clearvars -except k mrange l dates lengthdates lrange
end
end

```

REFERENCES

- Agnelli, G., Cacciani, A., & Fofi, M. 1975, *Solar Physics*, 44, 509
- Alibert, Y., Mousis, O., Mordasini, C., & Benz, W. 2005, *ApJ*, 626, L57
- Appourchaux, T. et al. 2014, *A&A*, 566, A20
- Arregi, J., Rojas, J. F., Sánchez-Lavega, A., & Morgado, A. 2006, *Journal of Geophysical Research: Planets*, 111
- Atkinson, D. H., Pollack, J. B., & Seiff, A. 1998, *J. Geophys. Res.*, 103, 22911
- Barrado-Izagirre, N., Rojas, J. F., Hueso, R., Sánchez-Lavega, A., Colas, F., Dauvergne, J. L., Peach, D., & Team, I. 2013, *Astronomy & Astrophysics*, 554, A74
- Bercovici, D., & Schubert, G. 1987, *Icarus*, 69, 557
- Bodenheimer, P., & Pollack, J. B. 1986, *Icarus*, 67, 391
- Bolton, S. J. et al. 2017, *Science*, 356, 821
- Boss, A. P. 1997, *Science*, 276, 1836
- Brasser, R., Morbidelli, A., Gomes, R., Tsiganis, K., & Levison, H. F. 2009, *A&A*, 507, 1053
- Brogi, M., de Kok, R. J., Albrecht, S., Snellen, I. A. G., Birkby, J. L., & Schwarz, H. 2016, *ApJ*, 817, 106
- Brouwers, M. G., & Ormel, C. W. 2020, *A&A*, 634, A15
- Cacciani, A., Dolci, M., Moretti, P. F., D'Alessio, F., Giuliani, C., Micolucci, E., & Cianno, A. D. 2001, *Astronomy & Astrophysics*, 372, 317
- Cacciani, A., & Fofi, M. 1978, *Sol. Phys.*, 59, 179

- Canny, J. 1986, *Ieee Transactions on Pattern Analysis and Machine Inligence*
- Chaplin, W. J. et al. 1996, *Sol. Phys.*, 168, 1
- Christensen-Dalsgaard, J. 2002, *Rev. Mod. Phys.*, 74, 1073
- Christensen-Dalsgaard, J., Duvall, T. L., J., Gough, D. O., Harvey, J. W., & Rhodes, E. J., J. 1985, *Nature*, 315, 378
- Civeit, T., Appourchaux, T., Lebreton, J. P., Luz, D., Courtin, R., Neiner, C., Witasse, O., & Gautier, D. 2005, *A&A*, 431, 1157
- Dick, D. J., & Shay, T. M. 1991, *Opt. Lett.*, 16, 867
- Domingo, V., Fleck, B., & Poland, A. I. 1995, *Sol. Phys.*, 162, 1
- Dowling, T. E. 1995, *Annual Review of Fluid Mechanics*, 27, 293–334
- Durante, D. et al. 2022, *Nature Communications*, 13, 4632
- Duvall, T. L., J., Dziembowski, W. A., Goode, P. R., Gough, D. O., Harvey, J. W., & Leibacher, J. W. 1984, *Nature*, 310, 22
- Duvall, T. L., J., Harvey, J. W., Jefferies, S. M., & Pomerantz, M. A. 1991, *ApJ*, 373, 308
- Edgington, S. G., & Spilker, L. J. 2016, *Nature Geoscience*, 9, 472
- Fegley, Bruce, J., & Lodders, K. 1994, *Icarus*, 110, 117
- Flasar, F. M. 1989, in *NASA Special Publication*, Vol. 494, 324–343
- Fletcher, L. N., Greathouse, T. K., Orton, G. S., Sinclair, J. A., Giles, R. S., Irwin, P. G. J., & Encrenaz, T. 2016
- Fletcher, L. N. et al. 2020a, *Planetary and Space Science*, 191, 105030
- Fletcher, L. N., Kaspi, Y., Guillot, T., & Showman, A. P. 2020b, *Space Science Reviews*,

- Folkner, W. M. et al. 2017, *Geophys. Res. Lett.*, 44, 4694
- Fuller, J. 2014, *Icarus*, 242, 283
- Fussell, J. A., Brazier, R. I., Davies, A. R., Isaak, G. R., McCleod, C. P., Morgan-Vandome, S. C., & Speake, C. C. 1995, 76, 452
- Gabriel, A. H. et al. 2002, *A&A*, 390, 1119
- Galperin, B., Sukoriansky, S., & Huang, H.-P. 2001, *Physics of Fluids*, 13, 1545–1548
- García-Melendo, E., & Sánchez-Lavega, A. 2001, *Icarus*, 152, 316
- Gaulme, P., Schmider, F.-X., Gay, J., Guillot, T., & Jacob, C. 2011, *Astronomy & Astrophysics*, 531, A104
- Gaulme, P. et al. 2008, *Astronomy & Astrophysics*, 490, 859, arXiv: 0802.1777
- Gaulme, P. et al. 2008, *A&A*, 490, 859
- Gaulme, P., Schmider, F.-X., & Gonçalves, I. 2018, *A&A*, 617, A41
- Gaulme, P., Schmider, F.-X., Widemann, T., Goncalves, I., Ariste, A. L., & Gelly, B. 2019, *Astronomy & Astrophysics*, 627, A82, arXiv: 1905.11078
- Gierasch, P. J., Conrath, B. J., & Magalhães, J. A. 1986, *Icarus*, 67, 456
- Giorgini, J. D. 2015, in *IAU General Assembly*, Vol. 29, 2256293
- Golitsyn, G. S. 1973, *Introduction to the dynamics of planetary atmospheres*.
- Gonçalves, I. et al. 2019, *Icarus*, 319, 795–811
- Grasset, O. et al. 2013, *Planet. Space Sci.*, 78, 1
- Grec, G., Fossat, E., & Pomerantz, M. 1980, *Nature*, 288, 541

- Guillot, T. 2005, *Annual Review of Earth and Planetary Sciences*, 33, 493
- Hart, A. B. 1954, *MNRAS*, 114, 17
- Harvey, J. W. et al. 1996, *Science*, 272, 1284
- Helled, R., Anderson, J. D., Podolak, M., & Schubert, G. 2010, *The Astrophysical Journal*, 726, 15
- Helled, R., & Stevenson, D. 2017, *ApJ*, 840, L4
- Helled, R. et al. 2022, *Icarus*, 378, 114937
- Hill, H. A., Stebbins, R. T., & Oleson, J. R. 1975, *The Astrophysical Journal*, 200, 484–498,
aDS Bibcode: 1975ApJ...200..484H
- Hubbard, W. B. 1999, *Icarus*, 137, 357–359
- Hueso, R. et al. 2017, *Geophys. Res. Lett.*, 44, 4669
- Iess, L. et al. 2018, *Nature*, 555, 220
- Ingersoll, A. P. et al. 2004, *Dynamics of Jupiter’s atmosphere*, Vol. 1, 105–128
- Ingersoll, A. P., & Pollard, D. 1982, *Icarus*, 52, 62–80
- Jesacher, A., Maurer, C., Schwaighofer, A., Bernet, S., & Ritsch-Marte, M. 2008, *Opt. Express*, 16, 2597
- Johansen, A., & Lambrechts, M. 2017, *Annual Review of Earth and Planetary Sciences*, 45, 359
- Johnson, P. E., Morales-Juberías, R., Simon, A., Gaulme, P., Wong, M. H., & Cosentino, R. G. 2018, *Planet. Space Sci.*, 155, 2
- Joye, W. A., & Mandel, E. 2003, in *Astronomical Society of the Pacific Conference Series*,

- Vol. 295, *Astronomical Data Analysis Software and Systems XII*, ed. H. E. Payne, R. I. Jedrzejewski, & R. N. Hook, 489
- Karkoschka, E. 1998, *Icarus*, 133, 134
- Kaspi, Y. et al. 2018, *Nature*, 555, 223
- Kong, D., Zhang, K., Schubert, G., & Anderson, J. D. 2018, *Proceedings of the National Academy of Sciences*, 115, 8499–8504
- Kramida, A., Yu. Ralchenko, Reader, J., & and NIST ASD Team. 2021, *NIST Atomic Spectra Database* (ver. 5.9), [Online]. Available: <https://physics.nist.gov/asd> [2016, January 31]. National Institute of Standards and Technology, Gaithersburg, MD.
- Lederer, S. M., Marley, M. S., Mosser, B., Maillard, J., Chanover, N. J., & Beebe, R. F. 1995, *Icarus*, 114, 269
- Leibacher, J. W., & Stein, R. F. 1971, *Astrophys. Lett.*, 7, 191
- Leighton, R. B., Noyes, R. W., & Simon, G. W. 1962, *ApJ*, 135, 474
- Lellouch, E., Goldstein, J. J., Rosenqvist, J., Bougher, S. W., & Paubert, G. 1994, *Icarus*, 110, 315
- Lellouch, E., & Witasse, O. 2008, *Planet. Space Sci.*, 56, 1317
- Lewis C. Roberts, J., & Neyman, C. R. 2002, *Publications of the Astronomical Society of the Pacific*, 114, 1260–1266
- Li, C., Lu, Y., & Zhang, S. 2021, *Optical Engineering*, 60, 073107
- Liu, S.-F., Hori, Y., Müller, S., Zheng, X., Helled, R., Lin, D., & Isella, A. 2019, *Nature*, 572, 355–357

- Lomb, N. R. 1976, *Ap&SS*, 39, 447
- Louden, T., & Wheatley, P. J. 2015, *ApJ*, 814, L24
- Luger, R., Agol, E., Foreman-Mackey, D., Fleming, D. P., Lustig-Yaeger, J., & Deitrick, R. 2019, *AJ*, 157, 64
- Luz, D. et al. 2006, *Journal of Geophysical Research (Planets)*, 111, E08S90
- Machado, P., Luz, D., Widemann, T., Lellouch, E., & Witasse, O. 2012, *Icarus*, 221, 248
- Machado, P., Widemann, T., Peralta, J., Gonçalves, R., Donati, J.-F., & Luz, D. 2017, *Icarus*, 285, 8
- Mankovich, C. R., & Fuller, J. 2021, *Nature Astronomy*, 5, 1103
- Markham, S., & Stevenson, D. 2018, *Icarus*, 306, 200–213
- Marley, M. S., & Ackerman, A. S. 2001, arXiv e-prints, astro
- Mayer, L., Quinn, T., Wadsley, J., & Stadel, J. 2002, *Science*, 298, 1756
- Meadowlark Optics, *Spatial Light Modulators User Manual*, x. X. P. S. 2022
- Militzer, B. et al. 2022, , 3, 185
- Mosser, B. 1995, *Astronomy and Astrophysics*, 293, 586
- Mosser, B., Maillard, J. P., & Mékarnia, D. 2000, *Icarus*, 144, 104
- Mosser, B., Mekarnia, D., Maillard, J. P., Gay, J., Gautier, D., & Delache, P. 1993, *A&A*, 267, 604
- Müller, S., Helled, R., & Cumming, A. 2020, *A&A*, 638, A121
- Pesnell, W. D., Thompson, B. J., & Chamberlin, P. C. 2012, *Sol. Phys.*, 275, 3
- Platt, B., & Shack, R. 2001, *Journal of Refractive Surgery*

- Pollack, J. B., Hubickyj, O., Bodenheimer, P., Lissauer, J. J., Podolak, M., & Greenzweig, Y. 1996, *Icarus*, 124, 62
- Provost, J., Mosser, B., & Chabrier, G. 1994, in *IAU Colloq. 147: The Equation of State in Astrophysics*, ed. G. Chabrier & E. Schatzman, 596
- Reach, W. T., Franz, B. A., & Weiland, J. L. 1997, *Icarus*, 127, 461–484
- Rolfe, E. J. 1988, *Seismology of the sun and sun-like stars. Proceedings.*
- Scargle, J. D. 1982, *ApJ*, 263, 835
- Schmider, F.-X., Fossat, E., & Mosser, B. 1991, *Astronomy and Astrophysics*, 248, 281
- Schmider, F. X. et al. 2007, *A&A*, 474, 1073
- Schmider, F.-X. et al. 2019, in *EPSC-DPS Joint Meeting 2019, Vol. 2019, EPSC–DPS2019–1613*
- Schou, J. et al. 1998, *ApJ*, 505, 390
- Serra, D. et al. 2019, *MNRAS*, 490, 766
- Shaw, C. L., GullEDGE, D. J., Swindle, R., Jefferies, S. M., & Murphy, N. 2022, *Frontiers in Astronomy and Space Sciences*, 9, 768452
- Showman, A. P., & Dowling, T. E. 2000, *Science*, 289, 1737–1740
- Simon, A. A., Wong, M. H., & Orton, G. S. 2015, *ApJ*, 812, 55
- Simon, M. C. 1986, *Applied Optics*, 25, 369–376
- Slinger, C., Cameron, C., & Stanley, M. 2005, *Computer*, 38, 46
- Stevenson, D. J. 2020, *Annual Review of Earth and Planetary Sciences*, 48, 465
- Sánchez-Lavega, A. et al. 2017, *Geophysical Research Letters*, 44, 4679–4686

- Tomczyk, S., Streander, K., Card, G., Elmore, D., Hull, H., & Cacciani, A. 1995, *Sol. Phys.*, 159, 1
- Tomczyk, S., Streander, K., Card, G., Elmore, D., Hull, H., & Cacciani, A. 1995, *Solar Physics*, 159, 1
- Toner, C., & Jefferies, S. 1993, *The Astrophysical Journal*, 415, 852
- Tsiganis, K., Gomes, R., Morbidelli, A., & Levison, H. F. 2005, *Nature*, 435, 459
- Ulrich, R. K. 1970, *ApJ*, 162, 993
- Valletta, C., & Helled, R. 2020, *ApJ*, 900, 133
- VanderPlas, J. T. 2018, *ApJS*, 236, 16
- Vorontsov, S. V., Zharkov, V. N., & Lubimov, V. M. 1976, *Icarus*, 27, 109–118
- Wahl, S. M. et al. 2017, *Geophysical Research Letters*, 44, 4649–4659, arXiv: 1707.01997
- Walsh, K. J., Morbidelli, A., Raymond, S. N., O’Brien, D. P., & Mandell, A. M. 2011, *Nature*, 475, 206
- West, R. A., Baines, K. H., Friedson, A. J., Banfield, D., Ragent, B., & Taylor, F. W. 2004, *Jupiter. The Planet, Satellites and Magnetosphere*, 79, ISBN: 9780521818087
- Widemann, T., Lellouch, E., & Donati, J.-F. 2008, *Planet. Space Sci.*, 56, 1320
- Wu, J., Liu, K., Sui, X., & Cao, L. 2021, *Opt. Lett.*, 46, 2908
- Wu, Y., & Lithwick, Y. 2019, *The Astrophysical Journal*, 881, 142