STOCHASTIC INVERSION FRAMEWORK FOR MONITORING EVOLVING SURFACE SHIP MASS PROPERTIES DURING ARCTIC OPERATION

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by Yolanda Chia-Yi Lin August 2018 © 2018 Yolanda Chia-Yi Lin ALL RIGHTS RESERVED

STOCHASTIC INVERSION FRAMEWORK FOR MONITORING EVOLVING SURFACE SHIP MASS PROPERTIES DURING ARCTIC OPERATION

Yolanda Chia-Yi Lin, Ph.D.

Cornell University 2018

The presence of Arctic sea ice currently limits surface ship travel in the Arctic for most of the year. However, due to rising temperatures from global climate change, Arctic waters are becoming increasingly navigable for a greater percentage of the year. As interest in surface travel within the Arctic increases in the coming years, the safety of a ship operating in this context must be considered. Specifically, the possibility of ice accumulation on the topside of a ship is heightened due to environmental factors within the Arctic, including the presence of sleet, snow, freezing rain, and freezing spray. This additional mass on exposed topside surfaces, at the most extreme, could result in capsizing of a vessel.

The present research develops a framework to monitor evolving mass properties for a ship in Arctic operation, in order to ensure safe travel through the Arctic. As with any real-world application, the data for this work are limited and noisy, and the system is complex. When the real-world data are limited, when the signals of interest are noisy, or when mechanistic models are unavailable, stochastic inference can enable informed decision making regarding the natural and built worlds. Thus, this work leverages stochastic inference in order to investigate the real-world problem of Arctic travel.

First, the dissertation presents a proof-of-concept for applying this framework to identify a single mass parameter for the Research Vessel (R/V) Melville with no icing and in quiescent seas, both at model-scale and full-scale. Second, the framework is extended to consider multiple mass parameters of the R/V Melville while undergoing potential ice build-up configurations. The third component of the dissertation looks outwards to the sea, to gauge the near-field wave forcing acting on the ship. In particular, the present work reports on a validation experiment of an existing algorithm to determine scale and sea state from an uncalibrated camera.

Taken together, the chapters prepare a foundation from which an ice monitoring system could be implemented. The chapters also provide insight to the specific challenges that exist for the full realization of the proposed framework. While the presence of ice is a focus of this work, the framework could easily be translated to a ship operating in within a context in which any mass property is evolving; those systems may require the monitoring of different mass parameters, but the underlying framework and approach proposed here would remain unchanged.

BIOGRAPHICAL SKETCH

Yolanda C. Lin grew up in Brea, CA. She graduated from Dartmouth College in Hanover, NH with an AB/BE in Engineering Sciences and a double major in Asian studies, which included a study-abroad in Tokyo, Japan. Yolanda then moved to Golden, CO to work at the National Renewable Energy Laboratory as an energy analyst. During this time, she also fell in love with the great outdoors and became a hiking, backpacking, and climbing enthusiast. Following a short but fascinating stay in the renewable energy industry, Yolanda returned to academia to pursue an MS in Structural Engineering, particularly in the topic of Earthquake Engineering, at the University of Colorado Boulder, with Professors Abbie Liel and Ross Corotis. Seeing the impact of both of her advisors' research and teaching, Yolanda was inspired to continue on for a PhD.

In the fall of 2014, Yolanda arrived at Cornell to begin her work with Professor Christopher Earls. In addition to her research, Yolanda was also dedicated to strengthening the community bonds within the School of Civil and Environmental Engineering (CEE), the College of Engineering, and the greater earthquake engineering community through her active leadership in the Earthquake Engineering Research Institute (EERI).

Within CEE, Yolanda is responsible for two main program developments during her time on the CEE Graduate Student Association (CEEGSA), from 2014-2018. As a first-year, she began the Graduate Student Seminar (GraSS), after seeing a need for students to practice giving presentations in a low-stakes environment. The seminar continues to benefit the graduate students of CEE. As CEEGSA President (2015-2016), she also doubled programming for the graduate student orientation. This new programming was aimed to ensure that cohorts knew each others' names and could begin building a sense of community at the very start of their doctorate or graduate journey, in addition to the existing programming that covered the academic logistics of graduate school.

Yolanda was also heavily involved with the Engineering Graduate Student Association (EGSA) within the College of Engineering from 2015-2018. After her experiences with CEEGSA, she felt that EGSA would be the appropriate vehicle to help support the student leaders within each of the department graduate student associations. Many graduate leaders face similar challenges, or have similar goals, and there was no way for these leaders to collaborate effectively, or even meet one another easily. In Fall 2016, she established the semesterly Engineering Graduate Leadership Dinner, where representatives from each GSA are invited to come together, share their successes and their challenges, and find opportunities for collaboration amongst their peers. As EGSA Vice President (2017-2018), she also created a structure to enable EGSA to provide funding for interdepartmental events.

Yolanda has also been an active member and leader within EERI. She has been the Cornell University representative to the EERI Student Leadership Council since 2015. In 2016-2017, she served as a EERI SLC Executive Committee member. In 2017-2018, she served with Daniela M. Martinez Lopez as EERI SLC Co-Presidents, overseeing the operation of the SLC and the annual Undergraduate Seismic Design Competition, an undergraduate design competition with over forty school teams (totaling over 300 undergradaute students) from around the world. Dedication

To my parents, and to my nephew, Kaden.

ACKNOWLEDGEMENTS

I would like to first thank my advisor, Christopher Earls, for the opportunity to pursue this research as a graduate student. I am grateful for the insight and knowledge he has shared with me, in research and in life. I would also like to thank my two committee members, Professors Derek Warner and Bill Philpot, for their helpful teaching, feedback, and input along the way.

This work would not have been possible without the support of the Office of Naval Research, through grant N000141612369, under the guidance and leadership of Paul Hess. I would also like to thank Sarah Punzi and Jayson Geiser for their support through experimental work, as well as Reis Seggebruch for his time as an undergraduate research assistant directly supporting my project. I would also like to acknowledge the Graduate School at Cornell University for their support through a Graduate Student Fellowship my first year, and for their continued programming and resources for graduate students during the course of their degree program. I also would like to thank the Civil and Environmental Engineering staff, to whom we all owe our thanks for ensuring that our research, classes, and programs move forward smoothly.

Before my time at Cornell, I was fortunate to have a rich history of strong mentorship. I would like to extend my continued thanks to Abbie Liel, Ross Corotis, Emily Newes, Steve Peterson, Vicki May, John Collier, and James Dorsey for their lasting and impactful roles in enabling, encouraging, and preparing me for the task of undertaking a PhD dissertation.

The Earls' Community in Hollister 411 is a special and unique group that I am thankful to have been a part of. Justyna, I am truly grateful we have marched through this experience together. Andrew, John, Wensi, Nick, and Vasilis, thank you for all your advice, company, and good humor.

I never would have made it to the finish line without an army of supporters in my family and friends. To my family – Mom, Dad, Richard, Christina, Kaden, and Sandy – thank you for always believing in me and rooting for me. To Tanya, thank you for your infectious enthusiasm and incredible skill in bringing people together, and for being a constant source of inspiration, integrity, and positivity. To Brianne, Frances, Marika, Jared, and Zimu, I am indebted to each one of you for playing formative, irreplaceable roles in my time in graduate school. Jen and Nozomi – thank you for welcoming me with open arms when I first arrived at this *other* college on a hill, and for continuing to be my Ithaca shepherds throughout these four years. To the many others in my Ithaca support network: thank you. Ithaca is home to an incredible number of talented, passionate, and amazing human beings who I am honored to call my friends.

I also am grateful also for the many friends who helped steady the ship from afar (or through visits into town!). A special thanks to Lauren, whose regular visits to Ithaca made me feel like a slice of Colorado had been teleported here. Sarah, thank you for being my rock and my cherished academic sister. Alden, Alice, Casey, Liz, Jen, David, Sean, and Jeff, I remain grateful and thankful for you continued friendship, guidance, and phone calls along the way.

|--|

| | Biog | graphic | al Sketch | iii | | |
|---|------|---|---|------|--|--|
| | Ded | ication | | v | | |
| | Ack | nowled | lgements | vi | | |
| | Con | tents . | | viii | | |
| | List | of Tabl | es | х | | |
| | List | of Figu | ures | xii | | |
| 1 | Intr | oductio | on | 1 | | |
| | 1.1 | Scope | and Organization | 2 | | |
| 2 | Stoc | Stochastic inversion for the roll gyradius second moment mass prop- | | | | |
| | ertv | in shir | os at full-scale and model-scale | 5 | | |
| | 2.1 | Intro | duction | 5 | | |
| | | 2.1.1 | Scope and organization | 8 | | |
| | 2.2 | Overv | view of modeling | 8 | | |
| | | 2.2.1 | Forward modeling | 9 | | |
| | | 2.2.2 | Inverse modeling | 10 | | |
| | | 2.2.3 | Computational framework | 13 | | |
| | 2.3 | Appli | cations | 17 | | |
| | | 2.3.1 | Inversion using noise-free simulated data ("inverse-crime") | 18 | | |
| | | 2.3.2 | Full-scale demonstration | 19 | | |
| | | 2.3.3 | Model-scale validation | 28 | | |
| | 2.4 | Discu | ssion | 33 | | |
| | | 2.4.1 | Full-scale discussion | 33 | | |
| | | 2.4.2 | Model-scale discussion | 33 | | |
| | 2.5 | Concl | usions | 35 | | |
| 3 | Mu | lti-para | meter stochastic inversion for first and second moment | | | |
| | mas | s prope | erties of model-scale ship with topside ice accumulation | 37 | | |
| | 3.1 | Introd | luction | 37 | | |
| | | 3.1.1 | Scope and organization | 40 | | |
| | 3.2 | Metho | ods | 41 | | |
| | | 3.2.1 | Forward Modeling | 42 | | |
| | | 3.2.2 | Inverse Modeling | 43 | | |
| | | 3.2.3 | Computational framework | 47 | | |
| | 3.3 | Appli | cation | 49 | | |
| | | 3.3.1 | Physical experiment | 50 | | |
| | | 3.3.2 | Single-parameter inversions | 54 | | |
| | | 3.3.3 | Multi-parameter inversions | 59 | | |
| | 3.4 | Result | ts | 64 | | |
| | | 3.4.1 | Prior forms | 67 | | |
| | | 3.4.2 | Prior support | 68 | | |

| | 3.5 | Discussion | 70 | | |
|----|--|--|-----|--|--|
| | 3.6 | Conclusions and recommendations | 75 | | |
| 4 | Validation experiment of a single-view image sequence algorithm to | | | | |
| | iden | tify scale and sea state characteristics | 77 | | |
| | 4.1 | Introduction | 77 | | |
| | | 4.1.1 Scope and organization | 79 | | |
| | 4.2 | Single-view image sequence algorithm | 80 | | |
| | | 4.2.1 Background, assumptions, and definitions | 80 | | |
| | | 4.2.2 Theoretical requirements and known limitations | 87 | | |
| | | 4.2.3 Source code verification | 89 | | |
| | 4.3 | Validation experiment | 94 | | |
| | | 4.3.1 Description of data | 94 | | |
| | | 4.3.2 Experimental results | 97 | | |
| | | 4.3.3 Experimental parameters | 100 | | |
| | | 4.3.4 Image processing parameters | 101 | | |
| | | 4.3.5 Algorithm input parameters | 103 | | |
| | | 4.3.6 Y-aware parameter selection | 106 | | |
| | | 4.3.7 Image processing parameters, <i>Y</i> -aware | 111 | | |
| | | 4.3.8 Sequence M | 114 | | |
| | | 4.3.9 Image processing parameters, Sequence M | 116 | | |
| | 4.4 | Discussion | 117 | | |
| | 4.5 | Conclusion | 119 | | |
| Α | Cha | pter 1 of Appendix | 121 | | |
| | A.1 | Markov chain Monte Carlo | 121 | | |
| | A.2 | Algorithms | 125 | | |
| D | Cha | ntar 2 of Annondiv | 120 | | |
| D | R 1 | Additional tables for icing configuration results | 130 | | |
| | D.1 B 2 | Additional icing configuration results by prior forms | 130 | | |
| | D.2 B 3 | Additional icing configuration results by prior support values | 134 | | |
| | B.4 | Additional 2-D icing configuration results | 136 | | |
| | D.1 | | 100 | | |
| С | Cha | pter 3 of Appendix | 140 | | |
| | C.1 | Convergence criterion | 140 | | |
| D | Cha | pter 4 of Appendix | 142 | | |
| | D.1 | Background | 142 | | |
| | D.2 | Use-cases and motivation for monitoring ocean waves | 142 | | |
| | D.3 | Ocean wave monitoring instrumentation | 144 | | |
| | D.4 | Ocean wave monitoring algorithms and/or post-processing | 150 | | |
| Bi | bliog | raphy | 152 | | |
| | 0 | 1 V | | | |

LIST OF TABLES

| 2.1 2.2 | Summary of model inputs and results for the simulated inversion. Summary of results to calculate SNR from single full scale exper- | 19 |
|------------|---|------------|
| 2.3 | iment | 22 27 |
| 2.4 | Summary of instrument uncertainties from model-scale experi- ment | 30 |
| 2.5 | Summary of full-scale from model-scale MCMC inversion results | 32 |
| 3.1 | Locations of added weights on Model 5748 (see Figure 3.5) | 52 |
| 3.2 3.3 | Quantity of added weights on Model 5748 for each configuration R/V Melville (Model 5748) ship particulars by icing configura- | 53 |
| 34 | tion, model-scale | 54 |
| 2.5 | tion, full-scale from model-scale | 55 |
| 3.5 | configuration, full-scale from model-scale | 58 |
| 3.6 | Prior support values by type and configuration, full-scale | 63 |
| 5.7 | sults for Icing Configuration 0 | 65 |
| 3.8 | Summary of full-scale from model-scale MCMC inversion re- | 66 |
| 3.9 | Summary of full-scale from model-scale MCMC inversion re- | 00 |
| 3.10 | sults for Icing Configuration 5 | 66 68 |
| 11 | Symbols used in the ocean free surface wave field video analysis | 80 |
| 4.2 | Summary of <i>synthetic WAFO sequence</i> wave field video attributes | 80 |
| 43 | for code verification | 91 93 |
| 4.4 | Video collection schedule for small-scale image sequences E, F, |)) |
| 4.5 | G, H, M, and N \ldots summary of ground truth scale β values | 95 97 |
| 4.6 | Summary of results for E1–E3, raw. The average value for E, | 21 |
| | raw, is used henceforth as the baseline comparison value for this wave field | 100 |
| 4.7 | Summary of results for varying experimental design parameters | 102 |
| 4.8 | Summary of results for video perspective correction, with and without image enhancement | 104 |
| 4.9 | Summary of sensitivity results for video E1 at incremental devi- | |
| 4.10 | ations from DME for slice <i>L</i> | 105 105 |
| | | - |

| 4.11 | Summary of results for video sequence E, ranging from 1 ob- served cycle to 10 observed cycles. The raw video sequence ob- |
|------|---|
| | served 20 cycles |
| 4.12 | Summary of results for video E1, E2, and E3, with <i>Y</i> -aware pa- |
| | rameter selection |
| 4.13 | Summary of results for Experimental parameters, <i>Y</i> -aware 110 |
| 4.14 | Summary of results for video perspective correction, along with |
| | image enhancement through post-processing, for <i>Y</i> -aware pa- |
| | rameter selection |
| 4.15 | Summary of results for video E1 at varying directions for slice <i>L</i> , |
| | labeled as difference in degrees from true direction of maximum |
| | energy (DME), for the <i>Y</i> -aware parameter selection |
| 4.16 | Summary of results for Sequence M, <i>Y</i> -aware |
| 4.17 | Summary of results image processing parameters, <i>Y</i> -aware 117 |
| B.1 | Summary of full-scale from model-scale MCMC inversion re- |
| | sults for Icing Configuration 2 |
| B.2 | Summary of full-scale from model-scale MCMC inversion re- |
| | sults for Icing Configuration 3 |
| B.3 | Summary of full-scale from model-scale MCMC inversion re- |
| | sults for Icing Configuration 6 |
| B.4 | Legend for Figures 3.10 – 3.13 |
| | |

LIST OF FIGURES

| 2.1 | Overview of Bayesian inference implementation through MCMC | 7 |
|---|---|--|
| 2.2 | Schematic of computational framework | 14 |
| 2.3 | "Inverse Crime" with SNR 1,500,000 | 20 |
| 2.4 | Full scale Melville roll data and time derivative | 21 |
| 2.5 | Schematic for calculating a distribution around T using instru- | |
| | ment precision for amplitude of a measured motion – in this case, | |
| | the amplitude of roll | 24 |
| 2.6 | Normal distribution vs simulated roll period distribution | 26 |
| 2.7 | Simulated roll period histogram plotted with normal distribu- | |
| | tion fit to the simulated roll period data | 27 |
| 2.8 | Full-Scale Trials – assuming measured ship displacement (cen- | |
| | ter), lower bound ship displacement (on right), and upper bound | |
| | ship displacement (on left) | 28 |
| 2.9 | Full-Scale from Model-Scale Trial 1 | 32 |
| 2.10 | The best-case full-scale inversion results (using lower-bound | |
| | displacement input parameters) compared to the design value | |
| | (dashed-red vertical line at 0.39B) | 34 |
| 2.11 | The full-scale from model-scale inversion results, in green, com- | |
| | pared to a numerically generated population based on the the ex- | |
| | perimental full-scale from model-scale k^{ron} uncertainty, in black. | |
| | The dashed-red vertical line at 0.383B is the experimental ex- | 0 - |
| | pected value for $k^{r, out}$ | 35 |
| 3.1 | Simple inversion framework representation. The seakeeping for- | |
| | ward model is SMP, in black. | 43 |
| 3.2 | Schematic of applied Metropolis-Hastings algorithm | 4 - |
| | Schematic of applica Metropolis Hastings algorithm | 47 |
| 3.3 | Schematic of the computational framework for the inverse mass | 47 |
| 3.3 | Schematic of the computational framework for the inverse mass property solution | 47 49 |
| 3.3 3.4 | Schematic of the computational framework for the inverse mass property solution | 47 49 |
| 3.3 3.4 | Schematic of the computational framework for the inverse mass property solution | 47 49 51 |
| 3.33.43.5 | Schematic of the computational framework for the inverse mass property solution | 47 49 51 51 |
| 3.33.43.53.6 | Schematic of the computational framework for the inverse mass property solution | 47 49 51 51 59 |
| 3.3 3.4 3.5 3.6 3.7 | Schematic of the computational framework for the inverse mass property solution | 47 49 51 51 59 60 |
| 3.3 3.4 3.5 3.6 3.7 3.8 | Schematic of the computational framework for the inverse mass property solution | 47 49 51 51 59 60 |
| 3.3 3.4 3.5 3.6 3.7 3.8 | Schematic of the computational framework for the inverse mass property solution | 47 49 51 51 59 60 62 |
| 3.3 3.4 3.5 3.6 3.7 3.8 3.9 | Schematic of dependent for the inverse mass property solution | 47 49 51 51 59 60 62 |
| 3.3 3.4 3.5 3.6 3.7 3.8 3.9 | Schematic of upplied filedopoils flability of goriality is the end of applied filedopoils flability of goriality is the end of the computational framework for the inverse mass property solution | 47 49 51 51 59 60 62 62 62 |
| 3.3 3.4 3.5 3.6 3.7 3.8 3.9 3.10 | Schematic of upplied filed opoils flability of goriality is the end of applied filed opoils flability of goriality is the end of the computational framework for the inverse mass property solution | 47 49 51 51 59 60 62 62 68 |
| 3.3 3.4 3.5 3.6 3.7 3.8 3.9 3.10 3.11 | Schematic of upplied incurpoint fractings algorithm $1 + 1 + 1 + 1 + 1 + 1 + 1 + 1 + 1 + 1 $ | 47 49 51 51 59 60 62 62 68 69 69 |
| 3.3 3.4 3.5 3.6 3.7 3.8 3.9 3.10 3.11 3.12 | Schematic of upplied hedopolis Habiligo algorithm $1.1.1.1.1$ Schematic of the computational framework for the inverse mass property solution | 47 49 51 51 59 60 62 62 62 68 69 69 69 |
| 3.3 3.4 3.5 3.6 3.7 3.8 3.9 3.10 3.11 3.12 3.13 | Schematic of upplied filebools filebools filebools upplied in the filebools for the inverse mass schematic of the computational framework for the inverse mass property solution | 47 49 51 51 59 60 62 62 62 62 63 69 70 |

| 3.15 | Icing Configuration 5 Results: By prior support | 71 |
|------|---|-----|
| 3.16 | Icing Configuration 0, Gaussian-Uninformative Form, Experi- | |
| | ment Prior Support | 73 |
| 3.17 | Icing Configuration 0, Gaussian-Uninformative Form, Un-Iced Prior Support | 73 |
| 3.18 | Icing Configuration 5. Gaussian-Uninformative Form, Experi- | |
| 0.10 | ment Prior Support | 74 |
| 3.19 | Icing Configuration 5, Gaussian-Uninformative Form, Un-Iced | |
| | Prior Support | 74 |
| 4.1 | Schematic of algorithm to compile the spatial frequency tempo- | |
| 4.0 | | 83 |
| 4.2 | Schematic of algorithm to compute the FFT of spatial frequency | 0.4 |
| 4.0 | charactic charactic terms in the second state of the second state | 84 |
| 4.5 | Schematic of algorithm to compute the spatio-temporal image. | 80 |
| 4.4 | temporal image | 87 |
| 15 | Schematic of algorithm to compute the average energy at each | 07 |
| 4.5 | temporal frequency in the image sequence | 87 |
| 46 | Simulated wave field used for the present verification evercise | 07 |
| 1.0 | generated with the WAFO toolbox for MATLAB with significant | |
| | wave height of 0.5m and a primary peak period of 2 seconds | 90 |
| 4.7 | (a) First frame of the <i>sunthetic</i> WAFO <i>sequence</i> wave field: (b) Di- | 20 |
| 1 | rection of maximum energy superimposed on 2D FFT of first | |
| | scene of the <i>synthetic WAFO sequene</i> wave field | 91 |
| 4.8 | (a) 2-D FFT slice compiled image from the synthetic WAFO se- | |
| | <i>quence;</i> (b) row FFT computed from each row of (a) and compiled | |
| | for each row, with resulting best-fit parabola | 92 |
| 4.9 | Maximizing α in order to solve for scale, β , for <i>synthetic</i> | 92 |
| 4.10 | (a) Spatio-temporal image of synthetic WAFO sequence; (b) row | |
| | FFT computed from each row of spatio-temporal image (a) | 93 |
| 4.11 | Averaged FFT (by column) of Figure 4.10b | 93 |
| 4.12 | The two grazing angles used to collect video data: on left, 15° | |
| | grazing angle; on right, 30° grazing angle | 95 |
| 4.13 | The two lighting conditions used to collect video data: on left, | |
| | ambient overhead lighting; on right, spotlight lighting | 96 |
| 4.14 | (a) First frame of <i>E1</i> video; (b) Direction of maximum energy | ~~~ |
| | superimposed on 2D FFT of first scene of E1 video | 98 |
| 4.15 | (a) Sequence E1 accumulated 2-D FFT slice composition; (b) row | |
| | FFI computed from (a) and compiled for each row, with result- | |
| | ing best-fit parabola for <i>Sequence E1</i> . Both images have dimen- | 00 |
| | Sions of L rows and N columns | 99 |

| 4.16 | (a) <i>Sequence</i> E1 spatio-temporal slice composition image at 30 fps: (b) row FET computed from (a) computed and compiled for |
|------|---|
| | all rows Both images have dimensions of L rows and N columns 99 |
| 4.17 | Averaged row FFT of Figure 4.16b |
| 4.18 | α search for Sequence E1 (top). E2 (middle), and E3 (bottom) 112 |
| 4 19 | Alpha search for Sequence M1 (top) M2 (middle) and M3 (bottom)116 |
| 4 20 | The FFT of the spatial-temporal image (Figure 4.15) filtered to |
| 1.20 | minimize vertical lines via masking the 2-D FFT |
| A.1 | Inverse model – building the Markov chain |
| B.1 | Icing Configuration 2 Results: By prior forms |
| B.2 | Icing Configuration 3 Results: By prior forms 133 |
| B.3 | Icing Configuration 6 Results: By prior forms 133 |
| B.4 | Icing Configuration 2 Results: By prior support 134 |
| B.5 | Icing Configuration 3 Results: By prior support 134 |
| B.6 | Icing Configuration 6 Results: By prior support 135 |
| B.7 | Icing Configuration 0, Gaussian-Gaussian Form, Experiment |
| | Prior Support |
| B.8 | Icing Configuration 0, Gaussian-Gaussian Form, Un-Iced Prior |
| | Support |
| B.9 | Icing Configuration 0, Uniform-Uniform Form, Experiment |
| | Prior Support |
| B.10 | Icing Configuration 0, Uniform-Uniform Form, Un-Iced Prior |
| | Support |
| B.11 | Icing Configuration 5, Gaussian-Gaussian Form, Experiment |
| | Prior Support |
| B.12 | Icing Configuration 5, Gaussian-Gaussian Form, Un-Iced Prior |
| | Support |
| B.13 | Icing Configuration 5, Uniform-Uniform Form, Experiment |
| | Prior Support |
| B.14 | Icing Configuration 5, Uniform-Uniform Form, Un-Iced Prior |
| | Support |
| C 1 | Convergence of MCMC chain of 1,000 hurn-in samples and |
| C.1 | 15 000 stored samples 141 |
| | |

CHAPTER 1 INTRODUCTION

The presence of Arctic sea ice currently limits surface ship travel in the Arctic for most of the year. However, due to rising temperatures from global climate change, Arctic waters are becoming increasingly navigable for a greater percentage of the year. As interest in surface travel within the Arctic increases in the coming years, the safety of a ship operating in this context must be considered. Specifically, the possibility of ice accumulation on the topside of a ship is heightened due to environmental factors within the Arctic, including the presence of sleet, snow, freezing rain, and freezing spray. This additional mass on exposed topside surfaces, at the most extreme, could result in capsizing of a vessel.

Given the possibility of ice accumulation while traveling within the Arctic, one strategy for retaining favorable seakeeping capabilities could be the prevention or removal of topside ice. Previous research has indicated that many of these strategies, including hydrophobic surfaces and physical removal, are financially infeasible, or physically ineffective beyond certain quantities of accumulation [59].

Thus, the present research is motivated by the the scenario in which ice has already accumulated beyond the reasonable limits of effective ice removal. A surface ship with significant topside ice accumulation will henceforth be referred to as an evolved ship configuration. In this context, a seakeeping prediction under the evolved configuration of the ship is desirable, in order to make an accurate assessment of seakeeping capabilities. Since direct measurement of an evolving vessel is not possible during operation, there is a need to understand current conditions of a vessel through indirect means. The present research develops a framework to monitor evolving mass properties for a ship in Arctic operation, in order to ensure safe travel through the Arctic. As with any real-world application, the data for this work are limited and noisy, and the system is complex. When the real-world data are limited, when the signals of interest are noisy, or when mechanistic models are unavailable, stochastic inference can enable informed decision making regarding the natural and built worlds. Thus, this work leverages stochastic inference in order to investigate the real-world problem of Arctic travel.

1.1 Scope and Organization

The following is an overview of the remaining three chapters of this dissertation. This dissertation is a collection of three papers, corresponding to Chapters Two through Four. Each chapter is an autonomous unit with an introduction and a conclusion. References for all chapters are collected at the end.

Taken together, the chapters prepare a foundation from which an ice monitoring system could be implemented. The chapters also provide insight to the specific challenges that exist for the full realization of the proposed framework. While the presence of ice is a focus of this work, the framework could easily be translated to a ship operating within a context in which any mass property is evolving; those systems may require the monitoring of different mass parameters, but the underlying framework and approach proposed here would remain unchanged.

Chapter Two: This study considers the problem of identifying a single mass parameter of a ship, one at 1:23 model-scale and one at full-scale [40]. The Research

Vessel (R/V) Melville is used for demonstration purposes and data availability, and the parameter of interest is the roll gyradius (k_{roll}). The roll gyradius is of primary interest because it is a quantity directly related to capsizing. This chapter details the specifics of the inverse formulation, computational framework, and input data that form the basis of the overarching method of this dissertation. The method is demonstrated with an ice-free 1:23 model performed in quiescent seas and zero velocity, as well as for the full-scale R/V Melville in following seas (a best-approximation for quiescent seas and zero velocity). Accelerometer data from the model-scale and IMU output from the full-scale are both used to demonstrate the capability and flexibility of this method. The chapter concludes with solutions from the Markov chain Monte Carlo simulations and a discussion of the results and discrepancies between the model- and full-scale. This work provides a proof of concept for the application of an inference framework to identify current mass parameters of a ship, and demonstrates how the method can leverage existing on-board instrumentation to accomplish this task.

Chapter Three: The third chapter builds upon the framework developed in Chapter Two. The work investigates a similar problem of identifying mass parameters of the R/V Melville [37]. However, this study extends the framework from Chapter Two in two ways. First, this study considers the additional presence of topside ice, whereas the first study considered only an un-iced ship configuration. Second, this work aims to uncover two mass parameters (roll gyradius and vertical center of gravity), rather than a single mass parameter (roll gyradius). The framework thus must be modified to handle the inversion for multiple parameters. This chapter also focuses on investigating the relevance of using particular Bayesian prior distributions. Discussion focuses on the trade-offs between certain prior distributions and the practicality of employing those

methods in a real-world setting.

Chapter Four: In a full implementation of the framework, the surrounding sea state would be a necessary input to the forward model within the framework. Chapters Two and Three focus on a ship in quiescent seas, at zero velocity, disturbed initially by an impulse force such that its subsequent unforced motion can be recorded. In contrast, Chapter Four examines an existing algorithm to interpret the surrounding sea state in order to close the loop for the proposed framework. That is, Chapters Two and Three focus inwards on the ship, while Chapter Four looks outwards to the sea, to gauge the near-field wave forcing acting on the ship. In particular, the chapter presents a validation experiment of an existing method to determine scale and sea state [38]. The algorithm of interest requires a single-perspective, uncalibrated video of the sea surface. The original study [70] relied on synthetic sea surface videos to demonstrate the method, and did not validate the method with real-world data. The present work contributes a necessary suite of validation experiments for this image processing method. Results, identified challenges, and future recommendations are offered at the conclusion of this chapter.

CHAPTER 2

STOCHASTIC INVERSION FOR THE ROLL GYRADIUS SECOND MOMENT MASS PROPERTY IN SHIPS AT FULL-SCALE AND MODEL-SCALE

2.1 Introduction

Ship stability considerations are typically treated during the initial design phase of a vessel, using the so-called *range of stability plot*, which is the hydrostatic restoring moment versus roll angle or righting arm versus roll angle in still water (Righting Moment (RM) curve or the Righting Arm (GZ) curve, respectively) [17] [77]. While this plot does not encompass all of the influences affecting the stability of a ship at sea (e.g., viscous roll damping, etc.) [75] [15], experience has shown that the consideration of this metric, along with the application of safety factors prescribed within rules from classification societies and other international standards, yields a design that performs satisfactorily in normal operational contexts [34]. However, the expectation of satisfactory performance, in this case, is largely experiential, and necessarily context dependent. A damaged ship, a ship operating in extreme weather, or one experiencing significant topside icing could not be expected to perform satisfactorily since then the initial design context would have been lost [34] [16] [33]. In such circumstances, making a seakeeping prediction regarding the actual condition of the ship, as well as a prognosis for the ship's future seakeeping performance in its new configuration, becomes desirable, especially in real-time (i.e., within a time interval that is consistent with the characteristic time scale of interest).

Understanding the current stability limits for the rolling degree of freedom

– the degree of freedom most important for predicting the propensity to capsize – is one dimension of ship performance where such real-time mass property information will be of great value [75]. As a *first step* towards being able to ascertain current ship roll properties within a seaway (i.e., excited by irregular seas) in real-time, the present research effort focuses on inferring actual roll properties of a given ship, at an arbitrary point in time (i.e., with cargo, equipment, topped-off fluids and fuel, etc.), through a consideration of its calm water roll period. The current paper describes work wherein a stochastic inverse problem is formulated and solved, so as to furnish an inference regarding current roll gyradius properties of the given ship at sea. An additional effort aimed at subsequent model-scale validation is also reported herein.

In both of these contexts (i.e., full-scale and model-scale), uncertainty in the roll gyradius inferences are quantified within a Bayesian context. This proposed approach leverages the shipboard inertial measurement unit (*IMU*) telemetry against motion predictions (made using seakeeping software, as part of the stochastic inverse problem framework). The interaction of these components leads to the quantification of uncertainty recording the actual second moment mass properties (in the current paper, attention is restricted to roll gyradius, but other mass properties could be similarly treated). Such second moment mass properties are of great importance in understanding the dynamic behavior of a vessel through the ship equation of motion.

Other researchers have studied a related inverse problems that aimed to use the ship itself, as a kind of wave buoy, so as to make inferences regarding the on-site wave spectrum of the irregular sea surrounding a given vessel [46] [45]. This previous work assumed that a linear transfer function (taking the wave



Figure 2.1: Overview of Bayesian inference implementation through MCMC

spectrum into a ship response spectrum) was known. Within such a context, shipboard *IMU* measured responses (e.g., heave, roll angle, and pitch angle) could be used as part of the inversion for the desired wave spectrum.

The present work differs from this prior research: the current goal is to make inferences regarding the actual roll gyradius (as opposed to inferring a wave spectrum) of a given vessel through the use of shipboard *IMU* measurements of the vessel calm water roll period. Ship motion predictions, obtained from well understood seakeeping software, along with the *IMU* telemetry, are subsequently incorporated into a *likelihood function* as part of a Bayesian setting for the stochastic inversion leading to a roll gyradius inference. Figure 2.1 provides a high level overview of how the Bayesian inference is implemented in this context.

2.1.1 Scope and organization

The current paper is organized as follows. Section 2.2 starts with an overview of modeling, including the forward model, inverse model, and computational framework that enables our method. Following that, Section 2.3 includes a description of three application test cases for the current framework. The first application, in Section 2.3.1, tests the integrity of the framework with noise-free simulated data. Second, the full-scale application (Section 2.3.2) covers the details of the testing program and results for the full-scale application. Model-scale validation follows the full-scale application (Section 2.3.3). The discussion in Section 4.4 analyzes the utility of the method and furnishes subsequent thoughts based on inversion results. Finally, conclusions emanating from the current study are provided in Section 2.5.

2.2 Overview of modeling

The current paper reports on stochastic inversion efforts. A second moment mass property for roll (i.e., the roll gyradius) is obtained from the consideration of a noisy measurement of the natural, still water (or near equivalent) roll period from a model-scale or full-scale vessel. This inversion is carried out in conjunction with a seakeeping software tool, acting as a "black box" within a Bayesian formulation of the problem.

For reference, the definition of the in-air natural roll period, T_n , as determined by pendulum theory is given by the following:

$$T_n = \frac{2\pi k_{roll}}{\sqrt{gGM_T}} \tag{2.1}$$

where k_{roll} is the roll radius of gyration, or gyradius, g is the local acceleration of gravity, and GM_T is the transverse metacentric height. It is pointed out that the proposed inversion strategy employs the different, but related response: the calm-water roll period.

2.2.1 Forward modeling

In order to implement the stochastic inversion, a forward model is a necessary component of the framework. A forward model takes inputs to deterministically calculate an output of interest. In the forward model scheme, the roll gyradius is an input and the output is the calm-water roll period.

The forward model in the current work is a well-known seakeeping software tool: the U.S. Navy Standard Ship Motion Program (SMP) [11]. This tool was originally intended to be a means to aid the incorporation of seakeeping considerations in hull design early in the design process. Provided an input file with the specifications of the vessel of interest, including the roll gyradius, SMP's output includes deterministic predictions of ship motion in regular waves, amongst many other responses. One of these responses includes the calm water roll period of the specified vessel. SMP is based on linearized strip theory [11]. For the validity of this theory, assumptions require that the ship length be much greater than the beam and draft. Also, ship motions are limited to small amplitudes, such that the hull section and waterplane properties correspond to calm water values. Known issues with the application of SMP include greater accuracy for computed pitch and heave motions than roll motions, as well as reduced accuracy for large amplitude motions where non-linear effects are no longer negligible, such as when the wave heights are greater than the draft. The SMP was validated in a previous version (SMP81, 1981) and, since the underlying theories remained intact, validation studies were not repeated for the current release (SMP95) [11].

The SMP95 forward model takes a detailed input file that describes the ship, including characteristics such as hull geometry and loading. In formulating an input file to use for SMP95, data from experimental results are used when available. In the absence of measured values, design values are used as a best approximation. In particular, care is given to matching the ship displacement, controlled by the draft line input, and the metacentric height (*GM*), which is affected by the ship displacement (which affects the metacenter, *KM*) and loading particulars, including the free surface correction (*FSC*) and the distance between the keel and the center of gravity (*KG*). Specifically, *GM* is equal to $KM - KG_{out}$, where KG_{out} is calculated in SMP95 as $KG_{in} + FSC$.

2.2.2 Inverse modeling

An inverse modeling process is implemented in order to determine a *posterior distribution* of an inertial mass property. In this case, the inertial mass property is the roll gyradius (k^{roll}). The known quantities include experimentally collected values for calm water roll period ground truth (T_{GT}) and signal to noise ratio (SNR). The latter of the two encapsulates measurement errors. In other words, given some observed roll period, T_{GT} , SMP95 models (instantiated with param-

eters such as k^{roll}) are applied to furnish plausible model instances that match our observed response (i.e., T_{GT}).

The inverse model relies on basic probability principles. Bayes' Theorem is as follows:

$$p(A|B) = \frac{p(B|A)p(A)}{p(B)}$$
 (2.2)

where *A* and *B* are events, p(A|B) is the conditional probability density of observing *A* given that *B* is true (known as the *posterior distribution*), and p(A) and p(B) are the independent probabilities of the occurrence of event *A* and event *B*, respectively.

With Bayes' Theorem in this application, k^{roll} is event A and the calm-water roll period T_{GT} is event B. The equation can be re-written as:

$$p(k^{roll}|T_{GT}) = \frac{p(T_{GT}|k^{roll})p_{prior}(k^{roll})}{p(T_{GT})}$$
(2.3)

where $p_{prior}(k^{roll})$ is the *prior distribution* that encodes prior beliefs and practical limits on k^{roll} , and $p(T_{GT})$ is a normalizing term such that the probability density will integrate to 1. The quantity $p(T_{GT}|k_n^{roll})$ is known as the *likelihood distribution*.

Furthermore, T_{GT} is taken to be noisy – uncertainty can stem from any number of sources, from sensor noise to experimental error. The noise is assumed to be both stationary and Gaussian with some mean, μ , and standard deviation, σ , as is consistent with sensor noise. In this case, T_{GT} is taken as μ , and experimental standard deviation for T_{GT} , denoted as σ_T , as σ (where μ and σ are standard nomenclature in the literature). The relationship between the observed ground truth event's value, or signal, and the noise, or uncertainty, can be captured in the SNR value. The SNR used in this study is calculated as the ratio of the mean divided by the standard deviation, or T_{GT}/σ_T .

In the proposed inversion framework, the SMP forward model closes the likelihood function (Equation 2.4):

$$p(T_{GT}|k^{roll}) = \frac{1}{\sqrt{2\pi\sigma_T^2}} e^{-\frac{1}{2\sigma_T^2}(T_{GT} - SMP(k^{roll}))^2}$$
(2.4)

where $p(T_{GT}|k^{roll})$ is the probability of the ground truth associated with the calm water roll period given a particular mass property (in this case, the roll gy-radius) T_{GT} is an experimentally collected ground truth calm water roll period value, σ_T is the standard deviation of the experimentally determined T_{GT} , and $SMP(k^{roll})$ is the output of the forward model, SMP95, for the calm water roll period, given the input model parameter k^{roll} .

Since the objective is to uncover the likely k^{roll} value given an observed roll period, one approach could be to maximize $p(T_{GT}|k^{roll})$ over all possible k^{roll} values in order to uncover the most probable k^{roll} . This, however, does not include the consideration of prior beliefs and experience concerning realistic k^{roll} values. Instead, a Bayesian approach is applied from Equation 2.3 in order to reflect an understanding concerning the uncertainty regarding k^{roll} . Using Bayes' Theorem, the probability distribution of k^{roll} will always be conditioned on our observed ground truth. The distribution for $p(k^{roll}|T_{GT})$ is built by the probability of the roll gyradius given the ground truth for the roll period from a sampling scheme. In support of our subsequent Bayesian inference, a *Markov chain Monte Carlo* (MCMC) sampling approach is implemented. This sampling scheme is chosen in order to effectively vary the pertinent model parameters within relevant model space. The result is a chain of values that provide an estimate for $p(k^{roll}|T_{GT})$. See Appendix A.1 for a detailed description of the method.

2.2.3 Computational framework

The computational framework adopted in the current work enables the implementation of an MCMC inversion with minimal dependencies on the specifics of a particular forward model. This allows for the use of one forward model, such as SMP95, but also builds in versatility for other ship motion predictive tools, suitable for future applications. The computational framework is implemented in the high-level computer language Python [55]. Python is chosen for its built-in flexibility in scanning large text documents and for its ease in communication with external programs to allow for a "plug and play" framework. A schematic of the overall computational inversion framework implemented in this work is presented in Figure 3.3. Mainly, the framework takes as input the expected ground truth period, T_{GT} , and its standard deviation, σ_T , denoted as the user inputs in the figure, and produces a histogram representing the distribution of k^{roll} , represented by histogram icon in the bottom right corner of the figure.

To begin the inversion, the inputs provided by the analyst are employed to initialize the necessary parameters within the computational framework. Experimentally obtained ground truth and experimentally calculated standard de-



Figure 2.2: Schematic of computational framework

viation represent a subset of such inputs. Measurement noise is assumed to be largely due to electronic sensor noise. As such, to assume a Gaussian form for noise is reasonable. This form of the measurement noise inspires a similar Gaussian form for our likelihood.

After the user inputs are read in and stored appropriately, the overarching function, <stoch_solve>, is called. This function is described in detail in Appendix A.2: Algorithm 1, and briefly outlined in the subsequent discussion. As inputs, <stoch_solve> takes the burn-in number, sample number, parameter prior support, variance, and expected ground truth values. This function returns the main objective of the inversion: a vector of numbers, stored in a Python list data structure with the name "data." This vector may subsequently be used in the generation of a histogram of the probable distribution of the parameter of interest, in this case, the inertial mass property roll gyradius. As a reminder, the framework is general: other model parameters (or collections of parameters) may be inverted for instead of just the roll gyradius.

As part of the stochastic inversion, <stoch_solve> calls a function, <burnin> (Appendix A.2: Algorithm 2). The <burn-in> function makes N_{burn} number of moves in the chain of values for our parameter of interest, but these results are not stored in the final sample chain. Instead, α is tuned and <burn-in> helps to ensure that the chain begins sampling in a high probability region of the posterior *probability density function* (PDF) for k^{roll} . The α parameter is tuned to reach a target 0.42 acceptance rate, as per Link [41].

To evaluate the plausibility of each instance of k^{roll} , the burn-in phase, where the function <gibbs> (described in Appendix A.2: Algorithm 4) is called. This function is used to loop over multiple model parameters. In this study, though only one parameter is being varied (*i.e.*, k^{roll}), the multi-parameter functionality exists for completeness and future applications of the method.

Within \langle gibbs \rangle , the Metropolis-Hastings function \langle MH \rangle (Appendix A.2: Algorithm 5) is called. This function is the heart of the MCMC inversion, and the main contents of Figure A.1. \langle MH \rangle begins by generating a candidate move within the chain of responses at step *n*. It does this by taking the previous, most recently accepted parametric candidate and imposing a small parametric perturbation, based on α , as described in Appendix A.1. After the candidate is generated, the candidate is categorized as admissible or inadmissible under the prior lower- and upper- bounds. If it is not admissible, the previous response is returned in place of a new candidate.

If the candidate k_n^{roll} is within acceptable limits, based on the specified support of the prior, the forward model is called in order to generate a new response value for natural roll period, T_n , given by Equation 2.1. The framework is designed to be modular such that the specific motion predicting software can be

replaced with minimal disruption to the rest of the MCMC inversion procedure, allowing for flexibility and versatility of this general framework to meet future needs. This modularity makes the forward model used here, i.e., SMP95, an effective "black box."

When SMP95 is called in the forward model evaluation phase, a Python script modifies a SMP95 baseline input file for the ship (*e.g.*, Melville.inp), so as to reflect the current MCMC-updated roll gyradius value at step n. The function <formod> (Appendix A.2: Algorithm 6) then opens and reads the output file from SMP95 in order to extract the calm water roll period (T_n , in seconds). After this new response is extracted, it is stored in a vector and the output file is closed. The new response is then returned to <MH> and used in the calculation of the likelihood function ratio, r_n .

Taking the inverse model discussion from Appendix A.1 and placing it into a computational setting: the forward model returns to \langle MH \rangle a value for T_n , or the calm water roll period at step n. The likelihood function of this candidate value, relative to the ground truth, is compared to the previous step, n-1, within the ratio of the likelihood functions, r_n . Though discussed in Appendix A.1, a few additional details are of note. If r_n is greater than or equal to 1, the candidate k_n^{roll} is accepted and returned as the "previous" response. If r_n is less than one, then a biased "coin toss" is performed, where r_n is compared to a uniform random number from 0 to 1. If r_n is greater than the generated number, k_n^{roll} is still accepted and returned as the "previous" response; if r_n is less than the generated number, k_n^{roll} is rejected as a candidate value. In the case of rejection, k_n^{roll} is overwritten with k_{n-1}^{roll} , and k_{n+1}^{roll} will be calculated off of the value k_n^{roll} = k_{n-1}^{roll} . This seemingly unintuitive procedure is of vital technical importance, ensuring detailed balance required for validity of the method [23].

The parameter value at step *n* is then passed back through \langle gibbs \rangle to \langle burn-in \rangle where α is modified based on whether the candidate at *n* was accepted or rejected in \langle MH \rangle . *n* is incremented by one and iterated until the number specified in \langle burn-in \rangle is reached and the burn-in phase is complete. At this point, the final α is passed back to function \langle stoch_solve \rangle , and \langle stoch_solve \rangle calls on \langle sample \rangle (Appendix A.2: Algorithm 3), which runs the samples from which values will be collected and returned for use in building up the posterior histogram. In the \langle sample \rangle phase, α is a static value, and a roll gyradius value is stored at every step *n*. Otherwise, both the \langle burn-in \rangle phase and the \langle sample \rangle phase follow the same algorithmic flow, incorporating the same family of functions, \langle gibbs \rangle , \langle MH \rangle , and \langle formod \rangle .

After the total number of samples has been run ($n = N_{burn} + N_{samp} = N_{total}$), the final chain of parameter values is printed into a text document for the user. At this point, one instance of the inversion is complete. The script can easily be run in a batch mode.

2.3 Applications

The proposed method is applied to the Research Vessel (R/V) Melville, for which full-scale and model-scale data are available, as a demonstration of the method. The R/V Melville is a retired Navy vessel that was operated by Scripps Institute of Oceanography, University of California. Overall, R/V Melville had a four-decade long history of serving as a vessel of science and research for ocean interests [44]. In this study, data are available for both the existing fullscale R/V Melville, at 85m long and weighing in at 3000 tonnes, as well as a 1:23 scale test-model scale version of R/V Melville. The model-scale testing was performed in the Maneuvering and Seakeeping (MASK) basin at the Naval Surface Warfare Center Carderock Division (NSWCCD) in West Bethesda, Maryland, while the full-scale R/V Melville data reported on herein was taken during a research campaign off the coast of California during September 2013. For reference, the full scale R/V Melville from model scale has a natural roll period of 10.49s, while the model-scale R/V Melville was measured to be 2.139 +/- 0.043s (+/- 2.0%) [53].

2.3.1 Inversion using noise-free simulated data ("inversecrime")

In order to build confidence in the inversion framework (i.e., to ensure correctness of algorithms and implementation in computer code), results are first presented from the MCMC inversion consisting of simulated data with no added noise as our ground truth; this is known as the "inverse crime" [21]. The simulated ground truth results are obtained from the same forward model, SMP95, as used when evaluating the likelihood function (Equation 2.4). The simulated k^{roll} result is fabricated with an assumed SNR value of 1,500,000 – high enough to proxy an "infinite" SNR. The prior for roll gyradius is assumed to be uniform within the lower bound of 0.2 and upper bound of 0.55, where values are normalized by the beam dimension of the ship and therefore unitless (standard practice for SMP95 input files). The uniform prior is implemented in order to

| Туре | Variable | Value |
|----------------------|----------------------------------|-----------|
| Forward model input | k ^{roll} input | 0.383B |
| Forward model output | Calm water roll period | 10.80 s |
| Inverse model input | SNR | 1,500,000 |
| Inverse model output | k^{roll} distribution mean | 0.383B |
| Inverse model output | k^{roll} distribution σ | 0.0003 |

Table 2.1: Summary of model inputs and results for the simulated inversion.

utilize a so-called *uninformative prior*, which allows the observed data to have greater influence over the posterior. If the inverse model can recover exactly the input that was used to run the forward model, then this builds confidence that the framework is running as intended. Inputs and results for this inversion are summarized in Table 2.1.

In Table 2.1, the distribution mean does indeed match the forward model input: both are equal to 0.383B. With such a large SNR value, the σ is expected to approach zero, but instead is equal to 0.0003B. The forward model outputs the calm water roll period only up to 2 decimal places, and this results in a 0.0004B window of k^{roll} inputs that can all result in the same T output. This 0.0004B range can be identified in Figure 2.3. Thus, due to this finite precision of the SMP95 output, the inverse model posterior output contains some variance in the distribution rather than the single spike that one would expect from an "infinite" SNR value.

2.3.2 Full-scale demonstration

After building confidence in the correctness of algorithms and coding, as furnished by results from Section 2.3.1, the framework can be applied to more



Figure 2.3: "Inverse Crime" with SNR 1,500,000

real-world contexts.

Description of testing program

Full-scale testing on the R/V Melville was conducted in 2013. For a detailed comparison study between model-scale and full-scale R/V Melville roll gyradius properties, the full-scale ship would be required to be in calm waters and subsequently excited to roll so as to replicate model-basin conditions. Since this full-scale test was conducted in open waters, this was not a condition that could be rigorously achieved. As a practical measure, the ship heading and speed were adjusted to be in following seas, where vessel speed matched the prevailing swell speed (i.e., so-called "wave riding").

R/V Melville had no rudder, but instead had steerable propulsor pods. As such, when the ship heading was oriented appropriately, the R/V Melville cap-


Figure 2.4: Full scale Melville roll data and time derivative

tain had the roll stabilizers deactivated and brought the ship to a forced wet roll period by modulating the propulsor angles at a rate that was consistent with the design roll period of 10.49 seconds. During this forced roll maneuver in following seas, R/V Melville achieved a peak roll angle of 12.31° from vertical. Following the maneuver, R/V Melville was allowed to oscillate until at rest, unforced. The first peak after the forcing ended was recorded at 9.16°. The resulting time history is presented in Figure 2.4. The maneuver was executed only once during the cruise.

Thirty seconds of data were collected from the on-board F150 *IMU* at a sampling frequency of 5 Hz and delivered in raw binary form [51]. A suitable Python script was written to generate a text file from these binary data in order that a fast Fourier transform could be performed. The subsequent roll period was found to be 10.4 seconds.

Roll precision for the *IMU* instrument is specified by the manufacturer as $\sigma = 0.025^{\circ}$, but the roll precision registered during calibration on the day of testing was recorded as 0.053°. The signal to noise ratio for the full-scale test

| Method | \hat{T}^{FS} [s] | $\sigma_{T^{FS}} \left[\mathbf{s} \right]$ | SNR ^{FS} |
|--------------------------|--------------------|---|-------------------|
| Interpolated Monte Carlo | 10.32 | 0.0805 | 128.18 |

Table 2.2: Summary of results to calculate SNR from single full scale experiment

was affected by the roll precision of the *IMU*, as well as the sampling rate at which roll angles are being collected.

Full-scale roll gyradius forward model inputs

The SMP95 forward model input file was modified from the design values to match the displacement on the day of the cruise. This displacement was observed at 2,421 + /-47.14 MT. For reference, the design value is 2,932.8 MT. Otherwise, all other forward model inputs are consistent with documented design values for the R/V Melville.

Full-scale roll gyradius inversion inputs

A mean value and SNR value for the full-scale test are needed to implement the MCMC method in order to infer a probability distribution for the k^{roll} . With only one time series, calculating a sample standard deviation is not trivial. The following is a discussion of the chosen method to calculate an effective SNR value to use, referred to as the Interpolated Monte Carlo, along with the corresponding results for a value for k^{roll} . Table 2.2 summarizes the subsequent results from the application of this method.

The present approach for reckoning the effective measurement SNR from the single time series data set, taken at full-scale, considers only the uncertainties

inherent in the *IMU* data collection process. Additional uncertainties, including damping and effects of changing the heading of the ship, are not considered due to limitations in the nature of the measured data.

Two challenges to arriving at an effective SNR value are (1) dealing with the *IMU* sampling rate of 5 Hz and (2) creating many plausible instances of the single time series, such that the distribution of measured values are consistent with instrument precision, and the needed statistics can be calculated from the population of these instances.

The first challenge is the sampling rate. The instrument samples a roll angle measurement every 0.2 seconds. Figure 2.5 is a schematic the unforced harmonic roll response of R/V Melville, along with how the proposed method for ascertaining reasonable SNR estimates is implemented. The true roll period, T = $t_2 - t_1$, is shown in Figure 2.5. However, t_2 and t_1 fall between the instrument sampling points, and must be solved for deterministically.

This can be calculated as:

$$t_1 = t_i - \frac{a_i \Delta t}{a_{i+1} - a_i}$$
(2.5)

and similarly, for t_2 :

$$t_2 = t_j - \frac{a_j \Delta t}{a_{j+1} - a_j}$$
(2.6)

Since $T = t_2 - t_1$, expressions from Equations 2.5 and 2.6 can be substituted to solve for T, resulting in:

$$T = t_j - t_i + \frac{a_i \Delta t}{a_{i+1} - a_i} - \frac{a_j \Delta t}{a_{j+1} - a_j}$$
(2.7)



Figure 2.5: Schematic for calculating a distribution around T using instrument precision for amplitude of a measured motion – in this case, the amplitude of roll

Returning, now, to the context of the existing experimental data, in order to measure one roll period, a reasonable choice to measure one full cycle is to use the peak-to-peak measurement of the experimental data (dashed line in Figure 2.4). Since the movement of the ship is sinusoidal, a linear interpolation can be used to deterministically calculate a time only around the zero-crossing in the data (i.e., the Taylor series expansion of the sine function is virtually linear within neighborhoods close to zero-crossings). This ability to interpolate a zero-crossing, however, is in conflict with the desire to measure peak-to-peak of the original data. Thus, rather than use the original data, the time derivative

of the position data (i.e., velocity) is applied to evaluate peak-to-peak period measurement. The numerical derivative is shown in Figure 2.4.

Now that sampling error due to time quantization has been accounted for, the second challenge is in arriving at an effective SNR measurement: creating many instances of the single time series in such a way that the instrument precision is preserved. Given a single experimental measurement – indicating a single time series from which to measure T^{FS} – each sampled point from the *IMU* is assumed to represent the mean value, along with an experimentally measured standard deviation, σ° , of 0.053°. As instrument noise, each noise instance is assumed to be independent and Gaussian (see 2.5).

Turning back to the schematic calculations, values at time t_i , t_{i+1} , t_j , and t_{j+1} all have a calibrated distribution around each point, $\sigma^{\circ} = 0.053^{\circ}$, representative of the known instrument precision on test day. This distribution around each recorded amplitude, at each time, can be combined with the deterministic linear interpolation to furnish a distribution of t_2 , t_1 , and ultimately T. Rather than the experimentally obtained amplitude values a_i , a_{i+1} , a_j , and a_{j+1} , Gaussian distributions – $N(a_i, \sigma^{\circ})$, $N(a_{i+1}, \sigma^{\circ})$, $N(a_j, \sigma^{\circ})$, and $N(a_{j+1}, \sigma^{\circ})$ – provide many instances of a_i , a_{i+1} , a_j , and a_{j+1} from which a distribution for T is constructed. Now the stochastic roll period, T, expression is:

$$T = t_j - t_i + \frac{N(a_i, \sigma^{\rm o})\Delta t}{N(a_{i+1}, \sigma^{\rm o}) - N(a_i, \sigma^{\rm o})} - \frac{N(a_j, \sigma^{\rm o})\Delta t}{N(a_{j+1}, \sigma^{\rm o}) - N(a_j, \sigma^{\rm o})}$$
(2.8)

Monte Carlo sampling is subsequently applied in order to generate a total of 10,000 instances of 2.8 roll periods. Relevant statistical values are calculated from the output of this simulation.

The resulting values from this simulation indicate a mean roll period value $\hat{T}^{FS} = 10.32$ seconds, with a standard deviation of 0.0805 seconds. These values result in an SNR value of 128.18. Figures 2.6 and 2.7 indicate that a normal distribution for *T* is a reasonable approximation for the majority of the values in the simulated roll period distribution.



Figure 2.6: Normal distribution vs simulated roll period distribution

Full-Scale Demonstration Inversion Results

With an input of SNR 128.18 and mean 10.32 seconds for the calm water roll period ground truth, one representative posterior histogram for the full-scale inversion is shown in Figure 2.8. Supporting results are in Table 2.3. Addi-



Figure 2.7: Simulated roll period histogram plotted with normal distribution fit to the simulated roll period data

| Trial | T_{GT} | k_{roll} Mean | $k_{roll}\sigma$ |
|---|----------|-----------------|------------------|
| Design | — | 0.39B | — |
| Trial 1: Measured displacement | 10.32 | 0.3654B | 0.003443B |
| Trial 2: Measured displacement, lower bound | 10.32 | 0.3696B | 0.003475B |
| Trial 3: Measured displacement, upper bound | 10.32 | 0.3608B | 0.003437B |

Table 2.3: Summary of full-scale MCMC inversion results

tional probability distributions from using displacements at the lower and upper bounds of the test day measurement can be found in the same figure, to provide a visual reference relative to the mean. A 6.3% difference exists between the full-scale result and the full-scale design value for k^{roll} when using the measured displacement. The lower-bound displacement results in a 5.2% difference, and the upper-bound displacement results in a difference of 7.4% difference. The "best case" scenario for the three displacements is achieved at the lower-bound of the displacement reading.



Figure 2.8: Full-Scale Trials – assuming measured ship displacement (center), lower bound ship displacement (on right), and upper bound ship displacement (on left)

2.3.3 Model-scale validation

As indicated previously, R/V Melville was selected for consideration in this work since both full-scale and model-scale roll data were available for this vessel. The current discussion now shifts to the model-scale validation of the proposed method.

Description of model testing

Model-scale experiment testing of the seakeeping performance for the R/V Melville was conducted during a three week period in March 2012 [44]. The testing was performed in the MASK basin at the NSWCCD.

The scale-model in this experiment was Model 5720, a 1:23 scale model of R/V Melville. The 1:23 configuration was chosen to allow for the testing of a targeted range of sea states in the MASK basin. Simplifications in the design of the model include: neglecting roll tanks, bow thrusters, topside superstructure detail, and propeller geometry.

The model-scale experimental test was conducted with the primary objective of providing time-synchronized model ship motion and wave measurement data for moderate speeds and unidirectional waves under both regular and irregular conditions. These data are used to validate the current inversion process and also aid in developing numerical ship motion prediction tools in the future. Properties reported from the testing include mass, vertical center of gravity, longitudinal center of gravity, transverse centers of gravity, moments of inertia for roll and pitch, and transverse metacentric height [44]. The resulting roll radius of gyration and calm water roll period values are of primary interest for the purposes of this framework study. The model-scale roll angle was measured in degrees and captured by a BF Goodrich vertical gyroscope (model VG34-0809-1). This is a miniature dual-axis electro-mechanical spinning mass gyroscope that provided the primary pitch and roll angle outputs.

To verify the load condition, an inclining test was performed with a Wyler inclinometer placed at the stern and on the centerline to record the roll angle. The Goodrich gryoscope was not used for this purpose because it provided insufficient accuracy for this application (Table 2.4). The design value of k_{roll} for the ship is 0.3900B (normalized by the beam dimension, B). For the model, a value of 0.3826B +/- 0.0069B (+/- 1.8%, $\sigma = 0.00345B$) for k^{roll} was achieved after ballasting.

| Sensor | Range [degrees] | Accuracy [degrees] |
|---------------------------------|-----------------|--------------------|
| Wyler Zerotronic | | |
| Type 3/3DK-13-097 | +/-10 | 0.0055 |
| Rosemount Goodrich | | |
| Vertical Gyroscope VG 34-0809-1 | +/- 60 | 1 |

Table 2.4: Summary of instrument uncertainties from model-scale experiment

Roll decay tests were performed at 0, 8, and 12 knots full-scale. For each speed, the carriage towed the model up to speed. When at speed, the tow line was slackened. The model was excited in roll by impulsively depressing either the port or starboard deck near midship. The model was then allowed to roll freely until the motion amplitude decayed to less than one degree. This process was repeated once for each side at each speed condition. Though this test was performed at three speeds, the current paper utilizes only the results from the roll decay tests moving at 0 knots.

Model-scale forward model inputs

For the model-scale based forward model inputs, parameters from [53] and [44] are used in place of design values where available. The free surface correction and vertical center of gravity were adjusted such that the metacentric height would match a height of 1.24m (modified from 1.09m). The hull lines were also modified to match a reported displacement of 2924 MT, updated from 2933 MT.

Model-scale inverse model inputs

From the data from both the pre-test roll decay test (using the Wyler instrument) and the in-test instrument (the Rosemount), the roll period was calculated along

with uncertainty parameters. Uncertainty analysis used in this report is conducted as per JCGM Guide to the Uncertainty in Measurement (GUM) [30]. The calm water roll period was found to have a value of 2.229 +/- 0.016 seconds at the model scale at 95% prediction limit (U95). The uncertainty value is calculated from the standard deviation of the roll decay test at 0 knots, as previously described. To convert the model scale tests to an equivalent full scale value, the roll period duration is divided by the value of $\frac{1}{\sqrt{scale}} = \frac{1}{\sqrt{23}}$. The U95 quantity is defined as:

$$U95 = \frac{t_{95}S}{\sqrt{1+\frac{1}{n}}}$$
(2.9)

where U95 = uncertainty with 95% prediction limit, S is the sample standard deviation, t_{95} is the inverse of the two-tailed student-t distribution with 95% confidence for 14 samples, and n is the sample size, 14. The equivalent full scale roll period is 10.69 +/- 0.0767 seconds. This results in σ equal to 0.03645s and an SNR value of 293. The mean will be used as the ground truth 10.69. For the inversion exercise, each MCMC chain is run such that the SNR is calculated with a standard deviation, σ , corresponding to the measured U95 prediction limit (*i.e.*, $U95 \approx 2\sigma$).

Model-scale roll gyradius inversion results

Three different MCMC sample chains corresponding to the single model-scale ground truth roll period were collected. Each MCMC inversion consists of a burn-in chain of 5,000 samples that are not retained as part of the final data set, and 15,000 subsequent samples that are collected and stored as part of the *posterior distribution* histogram construction. The noise standard deviation for the roll period used to drive the stochasticity in the subsequent MCMC sampling



Figure 2.9: Full-Scale from Model-Scale Trial 1

| Experiment | T_{GT} | k_{roll} Mean | $k_{roll}\sigma$ |
|-----------------------------|----------|------------------------|-------------------|
| Full-scale design | 10.32 | 0.39B | _ |
| Full-scale from model-scale | 10.69 | 0.383B | 0.0362 |
| | | | |
| Trial | T_{GT} | k _{roll} Mean | $k_{roll}\sigma$ |
| Trial 1 | 10.69 | 0.378B | 0.00158B |
| Trial 2 | 10.69 | 0.378B | 0.00192B |
| Trial 3 | 10.60 | 0.2700 | 0 00 2 00P |

Table 2.5: Summary of full-scale from model-scale MCMC inversion results

is taken as the measured model-scale value of 0.03645 seconds. Trial 1's result from the MCMC inversion is shown in Figure 2.9. The uncovered mean roll gyradius is 0.378B across all three sample chains, 1.3% error in comparison to the measured value of 0.383B, though 0.383B is at least two standard deviations away from the uncovered roll gyradius of any of the three inversion trials. At closest, in Trial 3, the measured value is $2.5\sigma_{trial3}$ from the inverted result.

2.4 Discussion

2.4.1 Full-scale discussion

This demonstration did not precisely uncover the design value k^{roll} , as shown in Figure 2.10. However, many of the influential SMP95 input values, capable of affecting the inversion, including the free surface correction and vertical center of gravity values, had to be approximated and assumed from design values as a best estimate, which could easily account for the current difference in values. Furthermore, since this is a full-scale demonstration, measurements for the true k^{roll} are unavailable for the day of the cruise, and the design value of 0.39B is used as a proxy which may or may not reflect the true k^{roll} value for the ship in the conditions of that particular day. Furthermore, since this is a design value, a representative distribution cannot be superimposed on the expected value since there is no data to support a likely standard deviation or uncertainty limit in the full-scale demonstration.

2.4.2 Model-scale discussion

When considering the carefully arrived at experimental uncertainty associated with the model-scale experiment, the proposed inversion scheme did not uncover the ground truth value for roll gyradius within two standard deviations for any of the three trials of the full-scale from model-scale inversions. The reason that this is surprising is that the accuracy of the computational formulation and implementation underpinning the stochastic inversion was verified by the "inverse crime" analysis in Section 2.3.1. In spite of the verified computational



Figure 2.10: The best-case full-scale inversion results (using lower-bound displacement input parameters) compared to the design value (dashed-red vertical line at 0.39B)

framework, the application of the method to the model-scale yields a k^{roll} value of 0.378B as compared with the measured, actual k^{roll} value of 0.383B, where 0.383B is outside of two standard deviations of the k^{roll} distribution. From this perspective, the results appear unsatisfying.

However, from a different vantage point, the results begin to look reasonable for this proof of concept. Figure 2.11 overlays the probability distribution from the full-scale from model-scale inversion results with a Monte Carlo generated population with mean 0.383B (the experimental expected value for k^{roll}) and standard deviation of 0.00345B: the precision on the ground truth measurement of k^{roll} , as measured experimentally. The expected value is demarcated by the vertical red dashed line at 0.383B. From this perspective, while the expected experimental value does not fall within two standard deviations of the inversion result distribution (based on the precision of the T_{GT}), the inversion result falls completely within the experimental distribution associated with the

ground truth k^{roll} .



Figure 2.11: The full-scale from model-scale inversion results, in green, compared to a numerically generated population based on the the experimental full-scale from model-scale k^{roll} uncertainty, in black. The dashed-red vertical line at 0.383B is the experimental expected value for k^{roll}

2.5 Conclusions

The proposed stochastic inversion results from the MCMC framework described herein are presented for a full-scale demonstration and a full-scale from model-scale validation study. The MCMC results for the roll gyradius (k^{roll}) are consistently low relative to the design value in both the full- and model-scale R/V Mellville measurements. The full-scale from model-scale ground truth calm water roll period is 10.69s in comparison to the demonstration full-scale value of 10.32s ($\sigma = 0.0805$) and a full-scale natural roll period design value of 10.49s. Curiously, the experimental full-scale value is shorter than the in air natural roll period, though the experimental full-scale, in water, value is expected to be longer. One contributing factor for this might be the difficulty in conducting full-scale field experiments, relative to the more controlled environment of model-scale tests. Additional field experiments would enhance the estimation of the full-scale roll period in order to provide the inversion scheme with a more confident ground truth value.

The obtained roll gyradius inversions varied with the two inversion demonstrations, as they are each furnished with different calm water roll period ground truths and SNR values. The full-scale demonstration inversion result of 0.365B with $\sigma = 0.0034B$ is compared to the design value of 0.39B, and the fullscale from model-scale experiment inversion result of 0.378 with $\sigma = 0.002$ is compared to the experimentally measured roll gyradius of 0.383B ($\sigma = 0.0034$).

Given the full-scale uncertainties and assumptions regarding important forward modeling parameters, the results from this study appear promising as a first-step to using a stochastic inverse problem framework in order to ascertain current ship properties using existing equipment and data on-board a ship of interest.

CHAPTER 3

MULTI-PARAMETER STOCHASTIC INVERSION FOR FIRST AND SECOND MOMENT MASS PROPERTIES OF MODEL-SCALE SHIP WITH TOPSIDE ICE ACCUMULATION

3.1 Introduction

While the presence of Arctic ice cover previously hampered shipping during most of the year, the increasingly navigable waters are becoming an active area of interest. The retreat of Arctic ice cover for a greater number of months in the year is projected to dramatically increase in the coming decades. For example, the Northern Sea Route is projected to increase from a two-week open water period in 2012 to a nine-week open water period in 2030 [19], with open water defined as that which contains less than 10% of sea ice coverage. Additionally, the Northwest Passage and Transpolar Route are both expected to change from 40% sea ice coverage year-round in 2012 to five-to-six weeks of open water by 2030 [19]. The increased navigability requires additional safeguards in order to ensure safe maritime operations within the Arctic.

Maritime operations within the Arctic present specific challenges from the point of view of seakeeping and stability. In the Arctic, the potential for ice build-up on exposed topside surfaces is increased due to heightened and prolonged exposure to cold, extreme weather. The formation of ice on a ship at sea is a phenomenon that has been studied extensively by Ryerson [58, 59]. There are a number of physical processes that may result in the accumulation of ice on a ship surface. The most prominent cause of ice accretion is the process of freezing spray. When the bow plunges into waves and swells, water is lofted over the

ship's bow. Bow slamming can result in a spray-entrained wind that becomes a cloud, and as the cloud is carried over the superstructure, freezing spray occurs [61, 58, 60, 59]. All of the foregoing icing mechanisms are of concern during Arctic maritime operations.

Given the possibility of ice accumulation while navigating through the Arctic, one strategy for ensuring favorable seakeeping properties and ship stability could be the prevention of ice from accumulating in a significant manner, thus eliminating the need to assess its presence and severity. Additionally, previous researchers have investigated options for ice removal on ship surfaces [59]. Noteable approaches include the following: hydrophobic surfaces, physical removal strategies, and use of instruments such as infrared emitters or lasers [59]. Many of these measures are financially infeasible, ineffective in reasonable time and quantity, and/or dangerous to personnel [59]. The details of ice prevention technologies are outside the scope of this present study. Instead, this work focuses instead on a scenario where ice has already accumulated beyond the reasonable limits of effective ice removal.

In spite of the unique environmental challenges within the Arctic, ship stability is typically considered in a nominal sense during the initial design phase of a vessel. A range of stability plot, either the hydrostatic restoring moment versus roll angle (righting moment, RM) or the righting arm versus roll angle in still water (righting arm, GZ), is typically used to assess ship stability [17, 77]. Experience has shown that a plot of this variety, along with applicable safety factors prescribed from classification societies and international standards, will result in a satisfactory design for normal operational contexts [34].

However, a ship operating in extreme weather, or a ship experiencing signif-

icant topside icing, can not be expected to perform to satisfaction since the initial design context has been lost [34, 16, 33]. In this scenario, making a seakeeping prediction under the actual condition of the ship in its new configuration (with evolved geometry and mass properties, such as a ship experiencing significant topside ice accumulation), becomes desirable. Since direct measurement of the evolving vessel (i.e., experiencing topside icing) is not possible, there is a need to understand the current condition of the vessel, during operations, using indirect means.

Other research has focused on a related inverse problem: in order to make inferences about the on-site wave spectrum in irregular seas, the ship itself is employed as a kind of wave buoy [46, 45]. The present work, in contrast, seeks to make inferences regarding the ship's evolving mass properties, rather than the surrounding sea spectra. In pursuit of this goal, the proposed approach employs ship motion predictions and inertial measurement unit (IMU) telemetry within a Bayesian setting for the stochastic inversion on a given mass parameter of interest (e.g., roll gyradius and vertical center of gravity).

The current work illustrates and extends an existing framework [40] to simultaneously assess the impact of icing conditions on two mass properties: the roll gyradius (k_{roll}) and the vertical center of gravity (VCG). The framework implemented in this study, previously verified and demonstrated in an initial study [40], leverages the existing shipboard IMU telemetry against motion predictions from existing seakeeping software as part of a stochastic inverse solution. In another previous study, each property was investigated independently, and each was quantified within a Bayesian context [39]. The current study, in contrast with the two previous aforementioned studies [40, 39], seeks to extend the work beyond independent single parameter inversions to now include multi-variable inverse solutions, employing only the same type of ground truth data available in the previous studies. Special attention is given to the prior parametric form (uniform, Gaussian, etc.) and prior support (the range of a random variable describing a mass property parameter) in order to determine a joint posterior distribution on the two variables of interest, given a measured vessel response from the IMU. We seek to determine an appropriate recommendation for prior parametric form and prior support for future operational use of the proposed framework.

3.1.1 Scope and organization

This paper first presents an overview of the methods used in mass property joint parametric inversion, including an overview of the forward modeling, inverse modeling, and computational framework implement in the study (Section 3.2). The following section (Section 3.3) focuses on the application of interest. Section 3.3.1 begins with an overview of the physical experiment (which utilizes 1:23 scale model of the R/V Melville) and resulting measurements. Section 3.3.2 describes the translation from model-scale experimental values to suitable inputs to the computational framework, in order to generate an inverse solution for one variable at a time, demonstrated for two *independent* parameters. An extension of the framework to two mass parameters, inverted simultaneously, is described in Section 3.3.3: integral to this discussion is the use of priors. Section 3.4 presents results from the multi-variable inversions, along with a discussion of those results in Section 3.5. Finally, conclusions and recommendations are offered in Section 3.6.

3.2 Methods

This work exploits stochastic inversion methods in order to arrive at plausible first- and second- moment mass properties of a surface ship, using only the input of a noisy motion measurement from the vessel's IMU. During the validation phase of the current work, the still-water roll period from a model-scale ship serves as this noisy motion measurement. We implement the proposed inversion framework to simultaneously invert for the roll gyradius (a secondmoment mass property) and the vertical center of gravity (a first-moment mass property). The inverse solution is expressed as a joint posterior (within a Bayesian context).

The inversion requires the repeated use of a seakeeping software tool, acting as the forward model (Section 3.2.1) and generally treated as a "black box" within a Bayesian formulation of this problem. The notional characterization of the seakeeping software as a "black box" emanates from the application of it as merely an input-output map, turning mass property parameters into predicted ship motions. Since the seakeeping prediction is treated abstractly, it enables the framework to be "plug and play," with the flexibility to accommodate whatever seakeeping predictions are required. The resulting mass property inference framework is represented in a simple schematic as shown in Figure 3.1, and in more detail, with reference to computational implementation, in Figure 3.3. Originally developed in a previous study [40], the stochastic mass property inversion framework is extended in the current work to uncover posterior distributions for two variables, rather than one. We examine in the current work a first-moment property in conjunction with a second-moment property.

3.2.1 Forward Modeling

The stochastic inversion framework requires the use of a "forward model," which, generally speaking, is a model that employs parametric inputs to deterministically calculate an output (or outputs) of interest (y). The forward model formulation in this case proceeds as follows: the input is the first- and/or second- moment mass property, and the output is the calm water roll period.

The forward model used in this work is the U.S. Navy Standard Ship Motion Program (SMP95) [11]. The tool is primarily used as a design aid to incorporate seakeeping considerations early in the hull design process. SMP95 requires an input file with the geometric and weight specifications of a vessel of interest. As an output, SMP furnishes deterministic predictions of ship motion in regular or irregular wave fields. One of the ship motion responses provided by SMP is the output of interest, the calm-water roll period (y = T).

The underlying formulation that is at the heart of the forward model, SMP95, is based on linearized strip theory [11]. Assumptions include: the ship length must be much greater than the beam and draft, and ship motions are limited to small amplitudes, such that the hull section and waterplane properties are comparable to calm-water values. SMP95 furnishes greater accuracy for pitch and heave motion predictions, as compared with roll motions [11]. Accuracy is also compromised for large amplitude motions where non-linear effects are no longer negligible, such as when wave heights exceed that of the draft. SMP was validated in a previous version, SMP81 (in 1981) [11]. However, since the underlying theories remain unchanged, validation studies have not been repeated for its current release, SMP95.

The input to SMP95, as mentioned above, requires preparation of a detailed input file to sufficiently describe the ship. Data to populate such a file are taken from experimental results when available; in the absence of physical model measurements, assumed design values for the given vessel are employed.

3.2.2 Inverse Modeling

For simplicity, the present section describes the proposed stochastic inversion framework in terms of scalar quantities. Vector quantities may be substituted without disrupting the approach to solution; thus, the scalar quantity is used without loss of generality.

In order to uncover the value (and associated uncertainty) of a mass property parameter, annotated here as θ , a stochastic inverse solution is implemented. A subsequent inverse solution is comprised of a mass property probability distribution, built through a sequence of many forward model simulations, shown in a simplified form in Figure 3.1.



Figure 3.1: Simple inversion framework representation. The seakeeping forward model is SMP, in black.

The current inverse solution exploits Bayes' Theorem:

$$p(\theta|y) = \frac{p(y|\theta)p(\theta)}{p(y)}$$
(3.1)

where $p(\theta|y)$ is the posterior mass property parameter distribution, θ represents the mass property parameter of interest, y is the observed ship motion, $p(y|\theta)$ is the conditional probability of observing the given ship motion at a particular mass property parametric instance (also known as the likelihood distribution), $p(\theta)$ is the prior distribution that encodes prior beliefs and physical limits on θ , and p(y) is a normalizing term to ensure that the probability distribution will integrate to 1. The observed ship motion y is obtained experimentally, either within a sea way or a model basin, and assumed to be noisy. The noise is approximated as independent, identically distributed, Gaussian, white noise so as to be consistent with electronic sensor noise within the IMU, with some mean, μ , and standard deviation, σ . The signal can further be characterized by its signalto-noise ratio (SNR), which captures the relationship between μ and σ . SNR is defined in this study as the ratio: μ/σ . For example, if the mean value is 10 and the standard deviation is 0.05, the SNR value is equal to 200.

The proposed stochastic inverse solution scheme, where a probability distribution function on θ is approximated, is an alternative to simply maximizing $p(y|\theta)$ over all possible θ in an effort to uncover the single most probable θ (e.g., maximum likelihood estimate, MLE). The Bayesian inverse solution approach is favored due to its ability to encode prior beliefs and to allow for expressiveness of a probability distribution.

To perform the described stochastic inversion, two input values are needed: 1) an experimentally observed ground truth measurement representing the ship's motion, such as the observed calm-water roll period; and 2) an understanding of the signal-to-noise ratio inherent in that measurement. Given these two inputs, the inverse solution framework will instantiate SMP95 models with proposed values for the mass property, θ . For a user-defined number of steps in the sequence, the inverse solution framework will subsequently accept or reject instances of θ , according to a specified acceptance criteria (here, the ratio of likelihood function as defined in Equation (2)), and the resulting collection of accepted values provide an estimate for $p(\theta|y)$ as a Markov chain of length *O* [23]:

$$f(\theta_n | \theta_{n-1}, \theta_{n-2}, ..., \theta_{n-O}) = f(\theta_n | \theta_{n-1}, \theta_{n-2}, ..., \theta_{n-O}, \theta_{n-O-1}, \theta_{n-O-2}, ...)$$
(3.2)

such that obtaining the current value at *n* is dependent only on the last O values, and not values beyond. In this case, the chain only depends on step n-1.

The ratio of likelihood function, *r*, is as follows in Equation 3.3:

$$r(\theta_{n^*})|\theta_{n-1}) = \frac{p(\theta_{n^*}|y)}{p(\theta_{n-1}|y)}$$
(3.3)

where θ_n is the proposed parameter value at step n, θ_{n-1} is the accepted value from the previous step, $p(\theta_{n*}|y)$ is the posterior probability of θ_n^* given observation, y, at step n, and $p(\theta_{n-1}|y)$ is the similar probability, but from the previous step, n-1. By Bayes' Theorem in conjunction with the forward model, the function, r, can be calculated as Equation 3.4, shown below with substitutions from Equation 3.1.

$$r(\theta_{n^*}|\theta_{n-1}) = \frac{\frac{p(y|\theta_{n^*})p_{prior}(\theta_{n^*})}{p(y)}}{\frac{p(y|\theta_{n-1})p_{prior}(\theta_{n-1})}{p(y)}}$$
(3.4)

Assuming that the distribution over y is stationary, the normalizing terms from the denominator of Equation 3.1 can be eliminated. The simplified equation, Equation 3.5, now reads:

$$r(\theta_{n^*}|\theta_{n-1}) = \frac{p(y|\theta_{n^*})p_{prior}(\theta_{n^*})}{p(y|\theta_{n-1})p_{prior}(\theta_{n-1})}$$
(3.5)

where $p(y|\theta_{n^*})$, the likelihood function, can be calculated from the forward model and ground truth values:

$$p(y|\theta_{n^*}) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(y-y_{n^*})^2}$$
(3.6)

where y_{n^*} is the response from the forward model instantiated on the proposed candidate mass property parameter θ_{n^*} , and σ is the standard deviation of the observed ground truth corresponding to the model input SNR (i.e., σ is the inverse of the sensor precision characterizing the instrument employed in measuring the ship motion). The probability distribution of $p(y|\theta_{n-1})$ takes a similar form, for the previously considered mass property instance at n - 1, instead of n^* . In this way, the ratio of likelihood functions may be calculated for each candidate parameter, θ_{n^*} , and consequently accepted or rejected based on the Metropolis-Hastings acceptance criteria [23]. According to the Metropolis-Hastings acceptance criteria, if r is greater than 1, the candidate is accepted at step n. If r is less than 1, then a "biased coin toss" is performed. Producing a uniformly distributed random number between 0 and 1, $U_n[0, 1]$, r is compared to this number: if r is greater than U_n , then θ_n is accepted, even though it is not "more likely" than the previous candidate; if r is less than U_n , then we reject θ_n .

Now that it is known how to compute, and how to utilize, the ratio in likelihood functions, the remaining component in the method hinges on generating candidates to be included in the posterior sampling. As alluded to above, after computing the ratio of the likelihood function and either accepting or rejecting the new value, we move forward by taking a small "step" away from the previously accepted value θ_{n-1} to a new value, θ_{n^*} .

$$\theta_{n^*} = \theta_{n-1} + \alpha * U[-1, 1] \tag{3.7}$$

where U[-1, 1] is a uniformly distributed random number between -1 to 1.

The result of repeated application of the foregoing is a chain of mass property parameter values that, when collected in bins, forms a histogram for $p(\theta|y)$. If the candidate θ_{n^*} is accepted, $\theta_n = \theta_{n^*}$ for this step in the sequence. If θ_{n^*} is not accepted, $\theta_n = \theta_{n-1}$ for this step in the sequence. This algorithm is described in detail in Figure 3.2. For a full description of the MCMC algorithm, as applied in this study, refer to Appendix A of [40].



Figure 3.2: Schematic of applied Metropolis-Hastings algorithm

3.2.3 Computational framework

The described inverse solution approach is realized as a flexible computational framework within the current work. For long-term feasibility and relevance,

this computational framework needs "plug-and-play" capabilities. Here we focus on the flexibility of the framework in order to accommodate the needs of "plug-and-play," meaning that the framework could be implemented on different systems with minimal disruption, especially with respect to components such as the forward model and ground truth input.

The framework is implemented in the high-level computer language Python [55], chosen for its built-in capabilities in scanning large text files and for its ease in interfacing with external, stand-alone programs. These capabilities allow the framework to achieve the desired "plug-and-play" flexibilty. A schematic of the overall computational inversion framework implemented in this work is presented in Figure 3.3. Each of the grey bars (and associated boxes) represent a function within the framework, with nesting indicative of the computational structure of how each function is implemented. Mainly, the framework takes as input the expected ground truth period, T_{GT} (the *y* described in the previous subsection), and its associated measurement noise standard deviation, σ_T , denoted as the user inputs in the figure, and produces a posterior distribution of the variables of interest, represented by the histogram icon in the bottom right corner of the figure.

The computational framework adopted in the current work enables the implementation of an MCMC inversion with minimal dependencies on the specifics of a particular forward model. As the schematic in Figure 3.3 suggests, the only dependency on the seakeeping software used is within the function, "formod." This builds in versatility for the use of nearly any forward model for future application. In other words, while we use SMP95 here, the framework can adopt a different seakeeping forward model with ease.



Figure 3.3: Schematic of the computational framework for the inverse mass property solution

3.3 Application

We are interested in detecting changes in the mass properties of a surface ship, particularly the roll gyradius (k_{roll}) and vertical center of gravity (VCG), measured relative to the keel. To do so, we consider six different configurations of ice build-up on the ship surface of the R/V Melville [54, 52, 59]. We collect physical experimental data from a model-scale version of the R/V Melville, and simulate the full-scale instance from that model-scale data of the R/V Melville using the forward model, SMP95. By harnessing both the physical experimental ground truth values and the simulated data from the forward model, the inversion framework generates statistical inferences on the two mass property parameters of interest. This work also establishes necessary experimental model-scale measurements and test configurations for use in formulating the prior distribution.

3.3.1 Physical experiment

This study uses experimental results from R/V Melville, David Taylor Model Basin (DTMB) Model 5748, a 1:23 scale physical model. Testing was conducted June 29, 2016 to July 7, 2016 at the Naval Surface Warfare Center, Carderock Division (NSWCCD) in Bethesda, Maryland. A previous model, R/V Melville, DTMB Model 5720, also 1:23 scale of R/V Melville, is reported on extensively for its ship particulars and seakeeping performance in [44, 53]. R/V Melville, DTMB Model 5748 is cast from the same mold as DTMB Model 5720 (Figure 3.4). Therefore, some parameters are assumed to remain consistent between the two models. Reporting on the physical experiment for R/V Melville, Model 5748 can be found in [54]. A summary of the experiment, along with relevant results, is included below in Tables 3.1 - 3.5. Six configurations of R/V Melville, Model DTMB 5748 are considered in this study. Here they are labeled Configuration 0, 1, 2, 3, 5, and 6. The different configurations are created by applying specific weights at each of the locations A-K, as shown in Figure 3.5. Configuration 0 represents the ship with no ice accumulation, shown in Figure 3.4, and is analogous to the previous study in [40]. Configurations 1, 2, 3, and 5 are representative of Overland Model icing scenarios as considered during the initial design of the actual R/V Melville [54].

Configuration 1 is based on one day of moderate ice accumulation on the starboard side only. Configuration 2 is based on two days of moderate icing mostly towards the bow of the ship. Configuration 3, in contrast to Configuration 2, is based on one day of moderate icing concentrated on the aft of the ship. Configuration 4 simulates on two days of heavy and extreme ice conditions, but the model was not tested due to its dangerously unstable properties.



Figure 3.4: Model 5748 tested at the Naval Surface Warfare Center, Carderock Division in June 2016 (Configuration 0, un-iced)



Figure 3.5: Locations of potential ice accumulation on Model 5748

Configuration 5 is generally symmetric with severe ice accumulation, and, like Configuration 4, also based on two days of heavy and extreme ice conditions.

| | Coordinates [m] | | | | | |
|----------|-----------------|--------|-------|--|--|--|
| Position | Х | Y | Ζ | | | |
| А | 1.281 | 0.000 | 0.356 | | | |
| В | 0.088 | 0.152 | 0.349 | | | |
| С | 0.883 | -0.152 | 0.349 | | | |
| D | 0.505 | 0.159 | 0.425 | | | |
| E | 0.505 | -0.159 | 0.425 | | | |
| F | 0.260 | 0.152 | 0.425 | | | |
| G | 0.260 | -0.152 | 0.425 | | | |
| Н | -0.715 | 0.158 | 0.311 | | | |
| Ι | -0.715 | -0.158 | 0.311 | | | |
| J | -1.330 | 0.154 | 0.197 | | | |
| Κ | -1.330 | -0.154 | 0.197 | | | |

Table 3.1: Locations of added weights on Model 5748 (see Figure 3.5)

Configuration 5 was adapted from its initial loading scenario due to testing feasibility, mostly by reducing the total load. Configuration 6 was created as an intentionally random distribution of ice mass, and based on the full-scale Melville deck plan. The Overland Model defines moderate icing as 0.7-2.0 cm/hour, and heavy/severe icing as 2.0 cm/hour or more [59].

The various topside icing scenarios considered were created by adding weights at 11 locations, labeled A to K. An isometric view of the ship is shown in Figure 3.5, with locations of positions A to K defined Table 3.1. The subsequent weight added at each location for each configuration is detailed in Table 3.2. Weight is added via 5 and 10 pound-force weights. Each weight, and all accompanying hardware, are carefully accounted for.

| | reactive weight [hoi] for each configuration | | | | | |
|----------|--|-------|-------|-------|-------|-------|
| Position | 0 | 1 | 2 | 3 | 5 | 6 |
| А | - | - | 5.00 | - | 10.00 | 10.00 |
| В | - | - | 10.35 | - | 5.35 | 5.35 |
| С | - | 5.34 | 10.34 | - | 5.34 | 5.34 |
| D | - | - | 15.37 | - | 10.38 | - |
| Е | - | 5.38 | 15.37 | - | 15.37 | 5.38 |
| F | - | - | 5.35 | - | 10.35 | 10.35 |
| G | - | 5.34 | 5.38 | - | 10.38 | 5.38 |
| Н | - | - | - | 15.35 | 5.35 | 10.35 |
| Ι | - | 5.33 | - | 15.34 | 5.34 | 10.34 |
| J | - | - | - | 10.35 | 10.35 | 5.35 |
| Κ | - | 10.36 | - | 10.36 | 10.36 | 15.35 |
| Total | 0 | 31.79 | 67.14 | 51.38 | 98.54 | 83.16 |

Added weight [lbf] for each configuration

Table 3.2: Quantity of added weights on Model 5748 for each configuration

For each icing configuration, the model is carefully weighed and ballasted in order to achieve target ship particulars, recorded in Table 3.3 for the model scale, and in Table 3.4 for full-scale from model-scale values. For full details on the model ballasting process, refer to [44, 54].

Each icing configuration is tested at rest for roll, pitch, and heave periods in still water. For each icing configuration, 30 roll periods, 5 pitch periods, and 5 heave periods are collected. Since this study utilizes only the roll results, the specifics of the pitch and heave motions will not be discussed further. To record calm-water roll period, the model is depressed about five degrees on the starboard side, held steady, then released. Motion from the first three complete roll periods is included in the subsequent analysis.

The ship's pitch, roll, and heave motions are recorded with a Kearfott KN-5050 inertial navigation system. The system includes a ring laser gyroscope and is capable of using a position-based input, such as GPS. The Kearfott naviga-

| Ship Particu | lars | Icing Configuration | | | | | |
|---------------------------|------------|---------------------|---------|---------|---------|---------|---------|
| Model-Sca | le | 0 | 1 | 2 | 3 | 5 | 6 |
| k _{roll} [in] | μ | 9.28 | 9.27 | 9.11 | 9.39 | 8.87 | 9.13 |
| | (σ) | (0.125) | (0.125) | (0.125) | (0.125) | (0.125) | (0.125) |
| k_{pitch} [in] | μ | 32.55 | 33.87 | 34.13 | 33.66 | 31.1 | 34.2 |
| | (σ) | (0.125) | (0.125) | (0.125) | (0.125) | (0.125) | (0.125) |
| VCG [in] | μ | 10.5 | 11.06 | 10.88 | 11.25 | 11.84 | 11.56 |
| | (σ) | (0.25) | (0.25) | (0.25) | (0.25) | (0.25) | (0.25) |
| Weight [lbf] | μ | 520 | 552 | 587 | 572 | 619 | 603 |
| | (σ) | (0.15) | (0.15) | (0.15) | (0.15) | (0.15) | (0.15) |
| $T\left[\mathbf{s} ight]$ | μ | 1.953 | 2.127 | 3.941 | 2.365 | 4.898 | 3.301 |
| | (σ) | (0.01) | (0.01) | (0.012) | (0.009) | (0.263) | (0.012) |
| $T \operatorname{SNR}$ | | 206 | 220 | 337 | 252 | 89 | 264 |

Table 3.3: R/V Melville (Model 5748) ship particulars by icing configuration, model-scale

tion system is rigidly mounted near the longitudinal center of gravity (via a mounting platform). The heading is accurate to 0.44 degrees, root mean square, while the roll and pitch angle is accurate to 0.03 degrees, root mean square [54]. Results from the roll motion test are summarized in Table 3. Values for the uncertainty measurement of the physical model-scale experiment are calculated by the method outlined in [35].

3.3.2 Single-parameter inversions

In order to utilize the experimental results from the physical test, model-scale values need to first be scaled up to full-scale for compatibility with SMP95 input requirements. Full-scale from model-scale ship particulars are summarized in Table 4.

As discussed previously, the user must define a roll period and signal-tonoise ratio for this measurement in order to perform the proposed stochastic

| Ship Particula | rs | Icing Configuration | | | | | |
|-----------------------------|------------|---------------------|-----------|-----------|-----------|-------------|-----------|
| Full-Scale | | 0 | 1 | 2 | 3 | 5 | 6 |
| <i>k_{roll}</i> [B] | μ | 0.384 | 0.384 | 0.377 | 0.389 | 0.367 | 0.378 |
| | (σ) | (0.0052) | (0.0052) | (0.0052) | (0.0052) | (0.0052) | (0.0052) |
| k_{pitch} [L] | μ | 0.246 | 0.256 | 0.258 | 0.254 | 0.235 | 0.258 |
| | (σ) | (9.45E-4) | (9.45E-4) | (9.45E-4) | (9.45E-4) | (9.45E-4) | (9.45E-4) |
| <i>VCG</i> [m] | μ | 6.136 | 6.464 | 6.358 | 6.575 | 6.767 | 6.756 |
| | (σ) | (0.146) | (0.146) | (0.146) | (0.146) | (0.146) | (0.146) |
| Displacement | μ | 2890 | 3061 | 3262 | 3175 | 3437 | 3352 |
| [MT*] | (σ) | (0.834) | (0.834) | (0.834) | (0.834) | (0.834) | (0.834) |
| $T\left[\mathbf{s} ight]$ | μ | 9.366 | 10.201 | 18.901 | 11.344 | 15.832 | 23.49 |
| | (σ) | (0.045) | (0.046) | (0.056) | (0.045) | $(0.06\ 0)$ | (0.263) |
| $T \operatorname{SNR}$ | | 206 | 219 | 337 | 252 | 264 | 89 |
| * MT in salt wa | ter | | | | | | |

Table 3.4: R/V Melville (Model 5748) ship particulars by icing configuration, full-scale from model-scale

inversion. The user must also define a specified number of "burn-in" samples and "stored" samples. The burn-in samples ensure that there is no dependence on the initial state during the stochastic inversion. Second, this burn-in period also tunes MCMC parameters that control the acceptance probability for θ [41]. For diagnostics purposes, the burn-in samples are saved within the computational framework, though they are not included when generating the posterior distribution histogram of θ . Stored samples are collected, and the binned samples approximate the posterior distribution of θ .

In choosing the number of "burn-in" and "stored" samples, ensuring that the MCMC sampling chain has sufficiently converged is critical, given the specific framework and model. A convergence criterion (Appendix C, [8]) can guide the choice of "burn-in" and "stored" sample values that the framework requires as a user-input. This builds confidence in the results of the generated posterior for a given mass parameter, θ .

Two sets of inversions were performed during preliminary work as part of the concurrent study. The first set inverts for the roll gyradius, k_{roll} ($\theta = k_{roll}$). The second set inverts for the vertical center of gravity ($\theta = VCG$). Each set is run independently; that is, when inverting for k_{roll} , VCG is deterministically assigned as the expectation (μ_{VCG}) of the experiment measurement, and vice versa. Both sets of inversions employ the experimentally determined roll period (μ_T) and associated SNR as the ground truth observation. That is: $y^* = \mu_T$, $SNR = \frac{\mu_T}{\sigma_T}$ = (mean of the experimental roll period measurements) / (standard deviation of the experimental roll period measurement).

Results

A summary of all the inversion solutions from the preliminary study is presented in Table 3.5. For a fully detailed presentation of these results, please refer to [39]. At present, a summary of the results and relevant implications of the preliminary study are presented for context for the multi-variable inverse solution, as well as for one of the prior supports described in Section 3.3.3.

From the preliminary results, it was observed that the roll gyradius alone is insufficient in characterizing the evolving mass properties in a ship experiencing ice accumulation. SMP95 assumes symmetry in the port-starboard (y-axis) direction. As such, it is expected that the inverse solution from Configurations 1, 5, and 6 would be most affected by this assumption. From the inverse solution, Configuration 1 seems to fare relatively well for both the posterior distributions of k_{roll} and VCG. This is likely due to its relatively minimal impact on achieved asymmetry under its lighter icing condition.
Configurations 5 and 6, however, appear greatly affected by this assumption. While the posterior distribution of VCG for both configurations aligns with expectations from the experimental uncertainty measurements, the roll gyradius posterior distribution solutions are not well supported by the experimental measurements, with relative errors of 78% and 28% for Configurations 5 and 6, respectively (as seen in Table 3.5). This can be explained by the relationship between k_{roll} and roll period, which, in SMP95 is positive and linear. In contrast, as VCG increases, the roll period increases quadratically, rather than linearly. The quadratic relationship between VCG and roll period is shown in Figure 3.6. The relationship shown is generated from the same base input file for Configuration 0, but instead varies one variable at a time while holding all else equal. VCG vs. roll period is plotted for $\mu + / - 3\sigma$. The red dashed lines represent the $\mu + / - 2\sigma$ (or U95) range for roll period given the μ , σ for the experimental values of VCG. Especially as the roll period exceeds about 12 seconds (or VCG exceeds 6.3 m), the effect of the quadratic VCG dominates over the linear k_{roll} . In comparison, the width of $\mu + / -2\sigma$ values of k_{roll} results in a roll period range of less than 0.5 seconds.

Each icing configuration is instantiated with deterministic, experimentally obtained measurements for pertinent input values (that is, mass parameter values not inverted for in the inverse solution). There is, however, still some discrepancy between 1) the forward model input file instantiated on some given mass property parameters and the associated forward model output versus 2) the experimentally obtained measurement for those same mass property inputs and the corresponding experimentally collected measurements for the roll period output. At a certain point, the disagreement between the experimental measurement and SMP95's output (e.g., forward modeling errors) for roll pe-

| Ship Particulars | | | Icing Configuration | | | | |
|-----------------------|------------|----------|---------------------|----------|----------|----------|---------|
| Full-Scale | | 0 | 1 | 2 | 3 | 5 | 6 |
| k _{roll} [B] | μ | 0.383 | 0.368 | 0.734 | 0.381 | 0.654 | 0.483 |
| | (σ) | (0.0023) | (0.002) | (0.0023) | (0.0018) | (0.0079) | (0.002) |
| % difference* | | 0.30% | 4.40% | 94.70% | 2.10% | 78.20% | 27.80% |
| VCG [m] | μ | 5.639 | 5.893 | 6.809 | 6.073 | 6.935 | 6.628 |
| | (σ) | (0.017) | (0.013) | (0.002) | (0.009) | (0.006) | (0.004) |
| % difference* | | 8.10% | 8.80% | 7.10% | 7.60% | 2.40% | 1.90% |

* Difference between physical model measurements and posterior distribution μ

Table 3.5: Results of R/V Melville (Model 5748) ship particulars by icing configuration, full-scale from model-scale

riod due to an unchanging VCG is too great for the roll gyradius to overcome, thus driving the roll gyradius to extreme values.

Even without considering the difference in experimental results and simulated results (model error), the range of roll periods that result from calculated uncertainty bounds on the experimental VCG and k_{roll} values are significant, and can also explain the behavior of these results. The k_{roll} U95 bound spans 0.43 seconds for the calm-water roll period. Meanwhile, in Figure 3.6, the VCGU95 bound spans 2.87 seconds, over six times greater than its k_{roll} counterpart. Given the experimental results, the roll period output is more sensitive to variations in VCG, and can easily account for the extreme values of some of the k_{roll} posterior distributions. The point in sharing these preliminary results is that interdependencies between desired mass property parameters influencing some measured ship motion response cannot be overlooked during application of the proposed framework. In such a case, the simultaneous, multi-parameter inversion capability of the proposed framework should be employed.



Figure 3.6: VCG relationship to roll period in SMP95 for Icing Configuration 0

3.3.3 Multi-parameter inversions

Section 3.3.2 illustrates the necessity to invert for the two related variables, simultaneously. With only one ground truth value, this is challenging to implement. In order to understand how to best incorporate available information on the ship, as well as the "evolved" topside ice distribution, this section focuses on the use of different prior forms (and supporting values) to instantiate appropriate prior distributions for use within the current Bayesian context (see Equation 3.5).

Prior forms

In previous studies [40, 39], the inverse solution was obtained using an uninformative prior on a single mass property parameter. The current work departs from previous work by investigating multiple mass parameters simultaneously, while also exploring the use of informative priors in this context. We investigate the application of informative priors on one or both of the mass parameters in question through the consideration of different prior forms: (1) a relatively narrow uniform prior (μ , +/ – 3 σ , with all other values rejected) and (2) a Gaussian prior (μ and σ in a Gaussian distribution). This makes for a total of four possibilities for use as prior distributions for our two-variable inversion: (1) U-U, uniform prior for both variables; (2) G-G, Gaussian prior for both variables; (3) Un-G, uninformative uniform prior (values rejected only outside those of physical or software limitation) for one variable, and a Gaussian prior for the other (k_{roll} and VCG, respectively); (4) G-Un, the same as (3) but in reverse (k_{roll} and VCG, respectively). These are depicted visually in Figure 3.7, where the prior form of the first variable, k_{roll} is in the left column, and the prior form of the second variable, VCG, is in the right column. Values for μ and σ are discussed in the following subsection.



Figure 3.7: Summary of sets of prior forms used in this two-variable study

Prior support

In addition to exploring various forms of prior probability density functions to apply to the current multi-variable inversion, we also consider the available values with which to apply as prior support (i.e., parametric range for which it is non-zero). Four prior supports are considered. First, we consider the use of mass property measurements taken from the un-iced model-scale experiment (labeled as *Un-Iced*). Second, the experimentally measured mass property values from each individual icing configuration are considered as prior support (labeled as *Experimental*). Third, using the mass property statistics obtained from the previous single-parameter study [39], we can apply these *Single-Parameter Inversion Informed* values as prior support for the two-variable inversion (labeled as *SPII*). Finally, we explore a hybrid of the single-inversion result for the expectation with the measurement uncertainty of the associated icing configuration experiment (labeled as *Hybrid*). A summary of the prior supports (1-4) for each configuration is shown in Table 3.6 and visually in Figures 3.8 and 3.9.



Figure 3.8: k_{roll} prior support values by type and configuration, full-scale. Mean values are plotted by the '.' marker and bars extend $+/-3\sigma$



Figure 3.9: VCG prior support values by type and configuration, full-scale. Mean values are plotted by the '.' marker and bars extend $+/-3\sigma$

| Mass Prop | erty | Icing Configuration Prior Support Values | | | | | | |
|--|------------|--|-----------|----------|----------|----------|----------|--|
| Paramet | er | 0 | 1 | 2 | 3 | 5 | 6 | |
| 1. Un-Iced Configuration (Design values for μ with Experimental σ) | | | | | | | | |
| <i>k_{roll}</i> [B] | μ | | | 0. | 39 | | | |
| | (σ) | | | (0.0) | 052) | | | |
| VCG [m] | μ | | | 5. | 89 | | | |
| | (σ) | | | (0.1 | .46) | | | |
| 2. Experin | nent n | neasureme | ents (Mod | el 5748) | | | | |
| k _{roll} [B] | μ | 0.384 | 0.384 | 0.377 | 0.389 | 0.367 | 0.378 | |
| | (σ) | (0.0052) | (0.0052) | (0.0052) | (0.0052) | (0.0052) | (0.0052) | |
| VCG [m] | μ | 6.136 | 6.464 | 6.358 | 6.575 | 6.767 | 6.756 | |
| | (σ) | (0.146) | (0.146) | (0.146) | (0.146) | (0.146) | (0.146) | |
| 3. SPII (SN | 1P95/ | 'Model 57 | 48) | | | | | |
| k _{roll} [B] | μ | 0.383 | 0.368 | 0.734 | 0.381 | 0.654 | 0.483 | |
| | (σ) | (0.0023) | (0.0020) | (0.0023) | (0.0018) | (0.0079) | (0.0020) | |
| VCG [m] | μ | 5.639 | 5.893 | 6.809 | 6.073 | 6.935 | 6.628 | |
| | (σ) | (0.017) | (0.013) | (0.002) | (0.009) | (0.006) | (0.004) | |
| 4. Hybrid | | | | | | | | |
| k _{roll} [B] | μ | 0.383 | 0.368 | 0.734 | 0.381 | 0.654 | 0.483 | |
| | (σ) | (0.0052) | (0.0052) | (0.0052) | (0.0052) | (0.0052) | (0.0052) | |
| VCG [m] | μ | 5.639 | 5.893 | 6.809 | 6.073 | 6.935 | 6.628 | |
| | (σ) | (0.146) | (0.146) | (0.146) | (0.146) | (0.146) | (0.146) | |

Table 3.6: Prior support values by type and configuration, full-scale

Framework updates

The currently employed framework was updated from its original form, as appeared in [40, 39], to allow for the simultaneous inversion of multiple parameters. While most of the framework components remain the same, a few updates have been included for this study. The Gibbs' algorithm is now implemented as Metropolis within Gibbs in order to handle the inversion for multiple mass parameters. At each step of the chain, this algorithm holds one variable steady while the other is slightly perturbed as if it would for a single variable [41]. The user-input was also updated to accept selection of prior form and specification of prior support, as needed.

3.4 Results

In total, 96 different icing mass property parameter inversions were considered, covering the four prior forms, four different prior supports, and six ice configurations. Key results are summarized in Tables 3.7 - 3.9. Some schemes could not be executed to completion, and are noted with a "-" in place of results. Their failure modes are discussed below. The pertinent results are also plotted in Figures 3.10 – 3.15, and remaining supporting tables and figures in Appendix B.1, Appendix B.2, and Appendix B.3. In Figures 3.10 – 3.15, the experimental μ values (full-scale from model-scale) for k_{roll} and VCG with a blue ".". Each blue "." is plotted within a grey rectangle, which indicates the $+/-3\sigma$ experimental values for each respective variable in each configuration. Figures 3.10 – 3.12 highlight the results of each form (Uniform-Uniform, Gaussian-Gaussian, Uninformative-Gaussian, or Gaussian-Uninformative, as indicated in each subplot title), considering all prior supports used with that specified form, with a "o" for the data points which correspond to the form of interest. Note that for each form, there should be four "o" data points. For example, for Icing Configuration 1 (Figure 3.11), if the form of interest is Uninformative-Gaussian, the "o" results will include the results using this form as applied to all four support options – essentially, all of column three Table 3.8 will be denoted as "o" marks, while the other three columns will be plotted with a "." for this form of interest, and so on for each of the subplots. Figures 3.13 – 3.15 highlight the results

| | Prior Type | | | | | |
|---------------|------------------------------|----------|---------|---------|----------|----------|
| Prior Support | Variable | | (1) U-U | (2) G-G | (3) Un-G | (4) G-Un |
| Un-Iced | <i>k_{roll}</i> [B] | μ | 0.3866 | 0.3881 | 0.3598 | 0.3900 |
| | | σ | 0.0113 | 0.0048 | 0.0170 | 0.0049 |
| | VCG [m] | μ | 5.6176 | 5.6099 | 5.8155 | 5.5914 |
| | | σ | 0.0899 | 0.0406 | 0.1227 | 0.0419 |
| Experiment | <i>k</i> _{roll} [B] | μ | 0.3770 | 0.3798 | 0.3251 | 0.3842 |
| | | σ | 0.0094 | 0.0047 | 0.0161 | 0.0054 |
| | VCG [m] | μ | 5.6927 | 5.6764 | 6.0492 | 5.6373 |
| | | σ | 0.0720 | 0.0381 | 0.1006 | 0.0455 |
| SPII | <i>k</i> _{roll} [B] | μ | 0.3834 | _ | - | 0.3828 |
| | | σ | 0.0049 | _ | _ | 0.0022 |
| | VCG [m] | μ | 5.6430 | _ | - | 5.6488 |
| | | σ | 0.0372 | - | - | 0.0225 |
| Hybrid | <i>k</i> _{roll} [B] | μ | 0.3840 | 0.3832 | 0.3873 | 0.3833 |
| | | σ | 0.0128 | 0.0050 | 0.0170 | 0.0053 |
| | VCG [m] | μ | 5.6363 | 5.6447 | 5.6108 | 5.6441 |
| | | σ | 0.1013 | 0.0413 | 0.1313 | 0.0443 |

Table 3.7: Summary of full-scale from model-scale MCMC inversion results for Icing Configuration 0

of each prior support (*Un-Iced, Experiment, SPII, or Hybrid,* as indicated in each subplot title) with a "o" for the data points which correspond to the support of interest, inclusive of all forms used with the specified support values.

Results for Icing Configuration 0, 1, and 5 are included in this section, as being relevant to the discussion and conclusions. Additional results for the remaining icing configurations are included in Appendix B.1 (for results formatted in a table by configuration), Appendix B.2 (for results by prior form), and Appendix B.3 (for results by prior support).

| | Prior Type | | | | | |
|---------------|------------------------------|----------|---------|---------|----------|----------|
| Prior Support | Variable | | (1) U-U | (2) G-G | (3) Un-G | (4) G-Un |
| Un-Iced | <i>k</i> _{roll} [B] | μ | 0.3893 | 0.3900 | 0.3632 | 0.3897 |
| | | σ | 0.0131 | 0.0054 | 0.0189 | 0.0053 |
| | VCG [m] | μ | 5.8553 | 5.8519 | 6.0161 | 5.8551 |
| | | σ | 0.0859 | 0.0382 | 0.1127 | 0.0372 |
| Experiment | <i>k</i> _{roll} [B] | μ | 0.3764 | 0.3794 | 0.2846 | 0.3845 |
| | | σ | 0.0065 | 0.0047 | 0.0304 | 0.0050 |
| | VCG [m] | μ | 5.9398 | 5.9248 | 6.4227 | 5.8881 |
| | | σ | 0.0399 | 0.0314 | 0.1421 | 0.0338 |
| SPII | <i>k_{roll}</i> [B] | μ | 0.3746 | _ | _ | 0.3680 |
| | | σ | 0.0011 | - | _ | 0.0020 |
| | VCG [m] | μ | 5.9364 | _ | _ | 5.9903 |
| | | σ | 0.0072 | - | - | 0.0172 |
| Hybrid | <i>k</i> _{roll} [B] | μ | 0.3667 | 0.3693 | 0.3887 | 0.3681 |
| | | σ | 0.0124 | 0.0047 | 0.0199 | 0.0050 |
| | VCG [m] | μ | 5.9970 | 5.9813 | 5.8573 | 5.9899 |
| | | σ | 0.0761 | 0.0303 | 0.1273 | 0.0318 |

Table 3.8: Summary of full-scale from model-scale MCMC inversion results for Icing Configuration 1

| | Prior Type | | | | | |
|---------------|------------------------------|----------|---------|---------|----------|----------|
| Prior Support | Variable | | (1) U-U | (2) G-G | (3) Un-G | (4) G-Un |
| Un-Iced | <i>k</i> _{roll} [B] | μ | 0.4107 | 0.3936 | 0.5829 | 0.3899 |
| | | σ | 0.0002 | 0.0050 | 0.0341 | 0.0047 |
| | VCG [m] | μ | 6.4736 | 6.5797 | 5.9535 | 6.5905 |
| | | σ | 0.0004 | 0.0138 | 0.1349 | 0.0130 |
| Experiment | <i>k</i> _{roll} [B] | μ | 0.3676 | 0.3672 | 0.4373 | 0.3677 |
| | | σ | 0.0122 | 0.0047 | 0.0282 | 0.0046 |
| | VCG [m] | μ | 6.6462 | 6.6475 | 6.4566 | 6.6462 |
| | | σ | 0.0303 | 0.0120 | 0.0835 | 0.0116 |
| SPII | <i>k</i> _{roll} [B] | μ | _ | _ | _ | 0.6450 |
| | | σ | _ | - | - | 0.0041 |
| | VCG [m] | μ | _ | _ | - | 5.6973 |
| | | σ | _ | _ | - | 0.0214 |
| Hybrid | k _{roll} [B] | μ | _ | 0.6449 | 0.3683 | 0.6470 |
| | | σ | _ | 0.0034 | 0.0587 | 0.0024 |
| | VCG [m] | μ | — | 5.7055 | 6.6331 | 5.6887 |
| | | σ | _ | 0.0186 | 0.1522 | 0.0158 |

Table 3.9: Summary of full-scale from model-scale MCMC inversion results for Icing Configuration 5

3.4.1 **Prior forms**

As expected, specifying the vertical center of gravity through a more "informed" prior causes significant errors in the roll gyradius. For example, for Configuration 0, from the *Experiment* prior support and *Uninformative-Gauss* forms, the proposed stochastic inversion framework yields a *VCG* that is 6.05 (1.5% difference from the experimental ground truth, the closest value for the *VCG* in Configuration 0); meanwhile, the roll gyradius for the same case is inferred to be 0.3251, a 15% error relative to the experimental ground truth. Again, for Configuration 0, the Un-Iced *Gaussian-Uninformative* combination recovered the k_{roll} exactly (0.39 for the expectation of both the prior support and the posterior), with a 5% difference relative to the prior support value of 5.89m. Relative to the experimental ground truth values, this was 1.5% difference in k_{roll} and an 8.9% difference for *VCG*. Results and errors are similar for Configuration 1, based on one of day of moderate ice accumulation on the starboard side only.

In the other configurations, these trends are exaggerated. When k_{roll} is bounded by more informative prior forms (uniform or Gaussian), the VCG tends to a larger value. Overall, this results in more accurate mass property inversion results, as k_{roll} is kept within a more plausible region of the parameter space and VCG has small-to-moderate errors. The obvious difference can be spotted by comparing the Uninformative-Gaussian support case to any of the other three support instances: relatively small errors of VCG are accompanied by exaggerated k_{roll} values, with k_{roll} errors well over 15%.

| Symbol | Meaning |
|----------------|--|
| Grey rectangle | k_{roll} and VCG ground truth values, range of $+/-3\sigma$ |
| | (based on exerimental full-scale from model-scale values) |
| blue "." | μ for k_{roll} and VCG ground truth values |
| | (based on exerimental full-scale from model-scale values) |
| black "o" | Inverse solutions of k_{roll} and VCG for category of interest |
| | (Category of interest listed in the plot title) |
| black "." | Inverse solution of k_{roll} and VCG for all other categories |
| | |

Table 3.10: Legend for Figures 3.10 - 3.15



Figure 3.10: Icing Configuration 0 Results: By prior forms

3.4.2 Prior support

Overall, the Un-Iced and Experiment prior supports yielded more plausible results during stochastic inversion. The *SPII* values pushed the k_{roll} values into regions of the mass property parameter space where SMP95 was wont to fail,



Figure 3.11: Icing Configuration 1 Results: By prior forms



Figure 3.12: Icing Configuration 5 Results: By prior forms

given the limitations of its basis in strip theory. These difficulties are increasingly obvious when the icing scenarios become more extreme (e.g., beyond Configurations 0 and 1).



Figure 3.13: Icing Configuration 0 Results: By prior support



Figure 3.14: Icing Configuration 1 Results: By prior support

3.5 Discussion

Generally, the Uniform-Uniform prior forms and the Gaussian-Uninformative prior forms were the most successful to employ during inversions, yielding the most plausible stochastic inversion results with respect to the experimental



Figure 3.15: Icing Configuration 5 Results: By prior support

ground truth measurements. The *Gaussian-Gaussian* and *Uninformative-Gaussian* prior forms were more restrictive. In the case of using a *Gaussian* prior form for the *VCG*, the same issues from the single-parameter inversion carried over; that is, the imposition of a more informative prior on the *VCG* forces the roll gyradius to assume implausible values, causing the inversion to fail.

The *Un-iced* and *Experiment* prior supports were most successful in yielding plausible mass property inversion results for each of the icing configurations. The *SPII* support was the most narrow, in terms of standard deviations, and at times the proposed stochastic inversion framework simply could not converge to a solution in a reasonable period of time.

The following discussion will thus include only the results of prior forms U-U, G-G, and G-Un, with prior support from *Un-Iced* and *Experiment*.

Errors from the two-variable study are greatly improved, as compared to the single parameter inversion [39]. The *Experiment* prior support is expected to re-

sult in a more accurate posterior distribution for the respective mass property parameter inversion results, relative to the *Un-Iced* support. This is true, particularly for the k_{roll} values, but the difference between results of the two prior supports at extreme icing (Icing Configuration 5) is moderate. For example, the difference between the experimental prior result and the ground truth using the G-Un form is 0.2% for k_{roll} and 1.8% for *VCG*. With the Un-Iced prior support, the difference is 6.2% for k_{roll} and 2.6% for *VCG*. This trend is generally continued across prior types and configurations. A 2-D histogram of Icing Configuration 0 comparing the results specifically between the *Experimental* prior support values and of the *Un-Iced* prior support values (across the *Gaussian-Uninformative* form) can be found in Figures 3.16 – 3.17, with additional results outside of this form and either of these supports in Appendix B.4. A relative color-scale for all of these figures is included above Figure 3.16.

Having an experimental prior ahead of time for every ice configuration possible is an unlikely and unrealistic expectation. However, what we see here is that the *Un-Iced* prior values can help provide sufficient informativeness to obtain plausible posterior distributions for stochastic inversions on multiple mass proprerty parameters. Furthermore, the less restrictive prior form (the *Gaussian-Uninformative* form) provides comparable results to the more informed supports. This enforces the previously posited observation that the roll period, *T*, is sensitive to changes in *VCG*. By keeping the roll gyradius in a more restricted probability space, the differences in *VCG* will not "push" the roll gyradius into extreme values, as was the case in the previous study [39].



Figure 3.16: Icing Configuration 0, Gaussian-Uninformative Form, Experiment Prior Support



Figure 3.17: Icing Configuration 0, Gaussian-Uninformative Form, Un-Iced Prior Support



Figure 3.18: Icing Configuration 5, Gaussian-Uninformative Form, Experiment Prior Support



Figure 3.19: Icing Configuration 5, Gaussian-Uninformative Form, Un-Iced Prior Support

3.6 Conclusions and recommendations

The accumulation of topside ice on a ship is a growing, critical concern in the Arctic, as waters become increasingly navigable and the risk of topside ice accumulation subsequently increases. This study extends an existing stochastic inversion framework to invert for two mass property parameters (rather than one), demonstrates the framework on four icing configurations in addition to an "Un-Iced" configuration, and makes recommendations for experimental measurements and real-world implementation of such a framework.

Improvement in results, as compared to the reported herein, could be made through the use of more sophisticated seakeeping prediction tools, so as to avoid the pitfalls accompanying the assumptions of linear strip theory. As stated in Section 3.2.1, SMP95 has known limited accuracy in considering the roll period. While the physical experiments also measured the pitch period and heave period, these measurements could not be utilized in this study due to limitations in SMP95. Furthermore, SMP95 is based on linear strip theory and also cannot consider mass asymmetry in the bow/aft direction (critical to the definition of realistic icing scenarios).

Working within these limitations in our available seakeeping forward model, we observe useful general trends in our inversion results. The implementation of an informative prior is recommended when inverting for multiple mass parameters of interest. Here, we show that a Gaussian prior form for the roll gyradius and an uninformative form for the *VCG* shows promise for these two parameter inversions. In practice, these findings suggest that future researchers could focus their physical data collection efforts on a careful under-

standing of the roll gyradius and associated uncertainty.

In the reality of balancing time, funding, precision, accuracy, and meaningful analysis, the results of this study can help inform a researcher seeking to prioritize these potentially competing factors. As for prior support values, an *Un-Iced* configuration is a practical recommendation for future applications of the proposed stochastic inversion framework (as in, a model-scale test of a ship of interest without ice building is a procedure that could be implemented as part of the protocol for this monitoring framework), while the availability of experimental mass parameters for every possible icing configuration for a given ship seems to be much less feasible. From an experimental perspective, the use of the *Un-Iced* configuration as the prior support is the most feasible option; from this study, we show that it is both feasible and reasonable. Though using the *Experiment* values for support is most accurate, the *Un-Iced* values resulted in posterior distributions for each variable that had errors only slightly larger than that of the *Experiment* support.

This study supports the use of a *Gaussian* distribution form for the roll gyradius, instantiated with values from an "Un-Iced" ship surface of a model-scale test, and an *Uninformative* form for the *VCG*, bound only by physical limitations. The implementation of a two-variable inversion has greatly improved results in attempt to uncover the mass parameters in a surface ship experiencing topside ice accretion.

CHAPTER 4 VALIDATION EXPERIMENT OF A SINGLE-VIEW IMAGE SEQUENCE ALGORITHM TO IDENTIFY SCALE AND SEA STATE CHARACTERISTICS

4.1 Introduction

In applied science and engineering contexts, there is a frequent desire to characterize the ocean free surface wave field. Motivation for such characterization can be found in such diverse applications as: weather monitoring [9], rogue wave warning [56], seakeeping predictions for surface ships [40], and even evasion of a pirate attack [62]. Various monitoring mechanisms for characterizing the ocean free surface have been developed, such as wave buoys [85, 36, 43, 6, 47, 29, 76, 12, 24], remote sensing instrumentation on satellites [36, 85, 22], Airborne Light Detection and Ranging (LIDAR) [28, 29], X-band radar [73, 27, 83, 78, 47], and various optics techniques [10, 66, 65, 63, 14, 72, 84, 69, 70].

The use of optics – specifically, the use of the visible wavelengths – in a sea state monitoring system is attractive due to its high spatial resolution and modest cost (i.e., employing off-the-shelf still cameras or video cameras). Additionally, with increasing computational capabilities, image processing and computer vision methods for this application have surged in the past two decades [65, 62, 3, 63, 70, 5]. However, due to the complications of obtaining well-characterized real-world datasets on sea states, optical wave imaging algorithms proposed within the literature are not always validated experimentally. At times, researchers have instead relied on synthetic datasets to substitute for well-characterized real-world data when assessing the performance of their al-

gorithms.

The scientific community relies on the process of verification and validation to build confidence in emerging scientific contributions [50, 48, 49, 57]. In order to formally discuss these concepts throughout the remainder of the present work, a few key terms are first defined to ensure consistency and clarity [49]. *Code verification* refers specifically to the process of determining the correctness of the numerical algorithms implemented as a software application. This process is intended to identify errors in the source code, but does not offer comparison between the real world and the simulated response. Source code verification is typically executed with a simple, well-known, and well-characterized test case, based on physical theory.

Validation is the process of determining the degree to which computational simulation (based on verified source code) is an accurate representation of the real world phenomena it is intended to represent. The comparison is quantified through the use of *validation metrics*, which measure the agreement between the computational and experimentally measured responses of interest. A *validation experiment* is composed of a suite or ensemble of many individual experiments in order to provide a critical evaluation of the *predictive* accuracy of a method. The evaluation may be considered on behalf of (1) the actual system in operational conditions that are less demanding than the anticipated use of the system, and/or (3) subsystems or individual components of the actual system [50, 48, 49, 57].

Finally, *sensitivity analysis* encompasses the process of determining how simulation results (i.e., responses of interest) depend on factors or parameters that instantiate the algorithm (i.e., inputs). Sensitivity analysis also looks at how outputs depend on explicit and implicit assumptions in the algorithm [49].

4.1.1 Scope and organization

A validation experiment of a single-view image sequence algorithm to identify scale and sea state characteristics is presented herein. This method, developed by Spencer et al. [69, 70], is novel for its combination of an uncalibrated single-view image sequence, rather than a stereo or multi-view image sequence, with the well-known dispersive properties of water waves, in order to recover scale and sea state. The algorithm can be applied using an inexpensive video camera (eg., GoPro Hero 5), making it an economical and accessible option for sea state observation. However, the algorithm has not been experimentally validated with reliable ground-truth data regarding the quantities the method seeks to recover, such as wavelength and dominant wave period.

The present paper first presents the algorithm of interest (Section 4.2), including relevant background, assumptions, and definitions (Section 4.2.1), as well as demonstrating the source code verification process in implementing the algorithm (Section 4.2.3). Section 4.3.1 describes the experimentally collected wave field videos for use in the present validation study. Section 4.3 also describes each of the validation exercises and accompanying results. Section 4.4 presents a discussion of the results from the overall validation experiment, and closing recommendations and conclusions are offered in Section 4.5.

| Symbol | Description | Definition | Units |
|-----------|--------------------|--------------------------------------|----------|
| λ | wavelength | horizontal distance (crest to crest) | m |
| H | wave height | vertical distance (trough to crest) | m |
| A | wave amplitude | vertical distance $H/2$ | m |
| δ | wave steepness | H/λ | m/m |
| T | wave period | time between crests | s |
| | - | passing a fixed spatial point | |
| С | wave speed | λ/T | m/s |
| f_s | spatial frequency | $1/\lambda$ | m^{-1} |
| f_t | temporal frequency | 1/T | s^{-1} |
| g | gravity of earth | 9.8 | m/s^2 |
| k | wavenumber | $2\pi/\lambda$ | m^{-1} |
| d | water depth | vertical distance (mean free | m |
| | - | surface level to bottom) | |
| | | | |

Table 4.1: Symbols used in the ocean free surface wave field video analysis

4.2 Single-view image sequence algorithm

The wave imaging algorithm of interest [69, 70] uses a single view camera image sequence, coupled with physics of deep water ocean waves, in order to extract frequency content and image scale without camera calibration. This section covers the background and assumptions of the imaging algorithm, and also describes and demonstrates the algorithm itself.

4.2.1 Background, assumptions, and definitions

The wave imaging algorithm of Spencer et al. [69, 70] adopts the common assumption that the ocean can be modeled as an ergodic, stationary random process [31], allowing the observation of a stationary wave in time and equating it to multiple realizations of a moving wave. Table 4.1 includes the relevant terms and definitions employed within the algorithm formulation. There are a number of parameters that characterize individual waves. The first value in the table, λ (meters), is defined as the wavelength, and can be measured as the horizontal distance from crest to crest. Wave Height, H, describes the vertical distance from trough to crest. Analytically, the upper limit of wave height can be derived as $H < \delta \lambda$, where $\delta \approx 0.143$ [31]. From experimental observations, the value for the upper limit of wave height has been recorded between $0.08 < \delta < 0.1$ [31]. Thus, if λ is known, H can be bounded by $H < 0.1\lambda$. Amplitude is equal to half of the wave height (A = H/2).

Another value in the table is the wave speed, *c*. The wave speed is defined as [68]:

$$c = \sqrt{\frac{g\lambda}{2\pi} tanh\left(\frac{2\pi d}{\lambda}\right)} = \sqrt{\frac{g\lambda}{2\pi} tanh(kd)}$$
(4.1)

Assuming that the ocean scene being imaged is in deep water (that is, the tanh(kd) term will approach 1 as the kd quantity goes to infinity). The deep water condition can be approximately satisfied when $kd \ge 2$. The equation can then be simplified to:

$$c = \sqrt{\frac{g\lambda}{2\pi}}.$$
(4.2)

Thus, applying the deep water assumption, the definition for c in Equation 4.2 will then be used in the following equations.

The Wave Period, *T*, using the above derived Equation 4.2 combined with the definition of $c = \lambda/T$, can be written as the following *dispersion relation*:

$$T = \frac{\lambda}{c} = \frac{\lambda}{\sqrt{\frac{g\lambda}{2\pi}}}.$$
(4.3)

The wavelength quantity still needs to be defined. From Equation 4.3, by squaring both sides, rearranging for λ , and substituting temporal frequency for the period ($T = 1/f_t$), an expression for λ can be arrived at:

$$\lambda = \frac{g}{2\pi} T^2 = \frac{g}{2\pi f_t^2}.$$
(4.4)

Process to determine scale

Spencer et al. [70] propose to determine scale with only an uncalibrated wave field video sequence. Though the reference [70] offers a more complete description of their algorithm, the following summary is furnished to enhance the stand-alone nature of the present paper. To begin, Spencer et al.'s definition of a scale factor, β , is repeated here (Equation 4.5).

$$\beta = \frac{\lambda_{pixels}}{\lambda_{meters}} \tag{4.5}$$

Substituting in the dispersion relations from Equations 4.3 and 4.4, β can also take the following form:

$$\beta = \frac{\lambda_{pixels}}{\frac{g}{2\pi}T^2} = \frac{2\pi\lambda_{pixels}}{gT^2}$$
(4.6)

where λ_{pixels} is defined as the number of pixels contained within the spatial dimension of one wavelength, as viewed from within a frame of the image capture sequence. Let x refer to the frame number, in time, ranging from x = 1 for the first frame and x = N for the final frame. For the first frame, the 2-D FFT is computed for the spatial image (Figure 4.1) in order to find the the *direction of maximum energy*, based on a search to maximize the average intensity over all possible directions. The direction of maximum energy is determined from the first frame only; that is, it is assumed that the direction of maximum energy does not change during the duration of the wave field video sequence (Figure 4.1,

red box around L_{SF1}). The first row of images in Figure 4.1 depicts a schematic of sequential spatial images, while the second row depends the corresponding 2D-FFT results of the spatial images. For each subsequent frame, the 2-D FFT is computed and a slice of the 2-D Fast Fourier Transform (FFT) is extracted in the identified direction of maximum energy (Figure 4.1, slices L_{SF2} , L_{SF3} , ... L_{SFN}). The slices are then sequentially added to a *spatial frequency temporal image*, where each slice in time occupies column x (Figure 4.1, bottom image). The resulting *compiled image* has dimensions of L by N.



Figure 4.1: Schematic of algorithm to compile the spatial frequency temporal image

The spatial frequency temporal image is then used to take the FFT of each

row (Figure 4.2, moving from the left image to the image on the right), in order to extract the temporal frequencies of waves with the wavelength $\lambda_{pixels} = L/y$. The *FFT of the spatial frequency temporal image* is then used with the following equations in order to determine image scale.



Figure 4.2: Schematic of algorithm to compute the FFT of spatial frequency temporal image

For each row y, the magnitude at a specified x coordinate corresponds to the energy of the wave with a given period, defined as:

$$T = \frac{N/v}{x} = \frac{N}{vx} \tag{4.7}$$

where N is the number of frames in the sequence and v is the frame rate, in frames per second. Combining Equations 4.6–4.7, the parabolic expression for y in the energy image is as follows:

$$y = \alpha x^2 \tag{4.8}$$

where α is equal to $\frac{2\pi Lv^2}{gN^2\beta}$. The value for α may be recovered through the optimization of the following objective function:

$$\alpha = \arg\max_{\alpha} \sum_{x} \operatorname{Energy}(x, \alpha x^{2})$$
(4.9)

where $\text{Energy}(x, \alpha x^2)$ is the pixel value at column x, row y of the 2-D FFT compiled image (i.e., the FFT amplitude value). Now, with a value for α , we can solve for β (Equation 4.10).

$$\beta = \frac{2\pi L v^2}{g N^2 \alpha} \tag{4.10}$$

Process to determine wave length

The second objective of the single-image sequence algorithm is to quantify the sea state. Again, the method utilizes a compiled image similar to the one described in the previous section. The same direction of maximum energy is used here to extract a single slice per frame (i.e., in time) from the spatial image and denoted as L_{SD} . The resulting *spatio-temporal compiled image* is comprised of one spatial slice of the grayscale image in the same dominant energy direction, L_{SD} , occupying column x, for a total of N columns. The y direction, again, has the same length L. The composite image has dimensions of L by N, as depicted in Figure 4.3.

Here, the focus is temporal frequencies, and so, the FFT of each row is computed (see Figure 4.4). If this is represented in an image, the row (y) is still the spatial location, but now the column x is the temporal frequency. Within this scenario, the value at (x, y) represents the energy of a given temporal frequency x at a given spatial location y.



Figure 4.3: Schematic of algorithm to compute the spatio-temporal image

Each column is then averaged so that there is one row vector, of size 1xN, representing the *average energy* (i.e, average FFT amplitude) at each temporal frequency in the image sequence (Figure 4.5).

At this point, the maximum value from the average energy plot may be calculated and the accompanying x value can be identified as the dominant temporal frequency. Alternatively, the data could also be fit to an appropriate ocean model to identify the dominant temporal frequency. The present study uses the first method to find the dominant temporal frequency, which immediately leads to the dominant wave period, T, from Equation 4.3, and wavelength λ , from Equation 4.4.



Figure 4.4: Schematic of algorithm to compute the row FFT of the spatiotemporal image



Figure 4.5: Schematic of algorithm to compute the average energy at each temporal frequency in the image sequence

4.2.2 Theoretical requirements and known limitations

The following discussion highlights some limitations that arise when applying the scale and sea state algorithm of Spencer et al. [70].

Spatial domain

According to the Nyquist–Shannon sampling theorem, the images must exhibit a scale of at least 2 times the wavelengths of interest in order to avoid spatial aliasing. For example, if a wavelength of 1m is of interest, there must be at least 2 pixels per one meter ($\beta = 2$).

Spencer et al. specify that there is no required field-of-view (FOV) to cover a minimum number of specified wavelengths, since the analysis is executed in the temporal domain. However, Spencer's study [70] also found, through consideration of synthetic data, that the algorithm was less reliable for scales $\beta > 40$. The inaccuracy was attributed to two factors: (1) fewer wavelengths per frame, and (2) the larger size of polygons generated in their synthetic model at scales greater than 40.

Temporal domain

Once again, applying the Nyquist-Shannon sampling theorem leads to the conclusion that the temporal sampling rate must be at least double the highest wave frequency within the wave field of interest in order to recover the frequency spectrum and avoid temporal aliasing. For example, if the highest frequency of interest is 1 Hz, then the sampling frequency (i.e., frame rate) must be at least 2 Hz. Also, for increased reliability, Spencer et al. suggest that at least 20 cycles of the dominant period should be captured within the video used for the algorithm [70]. Spencer et al. also tested the algorithm on synthetic sequences with fewer than 20 cycles. At a video duration of two dominant-periods (i.e., 4 seconds if the dominant period is 2 seconds), Spencer et al. determined that their algorithm recovered the ground truth dominant wave period within 16%.

4.2.3 Source code verification

The present section describes the approach taken to ensure that the computational implementation (i.e., source code) is an accurate representation of the algorithm proposed by Spencer et al.

Synthetic wave video

In order to *validate* the two single-perspective, uncalibrated wave imaging video algorithm, it is necessary to first *verify* the algorithm against "clean" synthetic data. The WAFO toolbox (Version 2.1) is a MATLAB-based software package for sea free surface modeling, including the simulation of linear and non-linear Gaussian waves that can be saved in video format [80]. It is maintained by the WAFO group at Lund University. The video used here is generated using the routine *seasim*. It is referred to here as the *synthetic WAFO sequence*, and denoted as "*WAFO*" in the subsequent tables. This video is generated from the JON-SWAP spectrum, a popular ocean wave spectrum suitable for simulating deep water waves [80]. The sea surface scene that is generated from the WAFO toolbox for this verification exercise has a dominant wave period of 2 seconds, and a significant wave height of 0.5 meters. The video is scaled at 6.70 pixels/meter.

Based on the Nyquist-Shannon theorem, the foregoing, prescribed sampling frequency should be sufficient to capture the frequency content of this video. In this case, the wave frequency is 0.5 Hz, and thus, the minimum sampling



Figure 4.6: Simulated wave field used for the present verification exercise, generated with the WAFO toolbox for MATLAB, with significant wave height of 0.5m and a primary peak period of 2 seconds

frequency is 1 Hz; this is far less than the 10 Hz sampling rate employed in this verification work.

The video perspective is directly overhead, and thus does not require the consideration of a perspective, or tilt, correction. Perspective correction is the application of an affine or projective transformation in order to rectify any distortion in original image in order to yield a computed, equivalent overhead view. The analyzed frame is 258 by 258 pixels, and 500 frames are considered (50 seconds of video, and 25 wave cycles captured). Given the scale, β , of 6.7 pixels per meter, the wavelength of 6.22 meters, and the pixel dimensions of the analyzed frame (258 pixels), it is expected that each slice *L* captures 6 spatial wavelengths. Table 4.2 summarizes the specific characteristics of the *synthetic WAFO sequence* employed in the source code verification.

Synthetic wave video results

Figures 4.7 - 4.11 show a demonstration of the single camera image sequence method. The first frame of the sequence is shown in Figure 4.7(a). This 2-D

| Characteristic | Value | Units |
|-----------------------------|-----------|-------------------|
| λ | 6.22 | [m] |
| T | 2 s | [s] |
| f_t | 0.5 s | [Hz] |
| eta | 6.7 | [pix/m] |
| С | 3.12 | [m/s] |
| Analysis window | 258 x 258 | [pixels x pixels] |
| Frames N | 500 | [frames] |
| Temporal sampling frequency | 10 | [frames/s] |

 Table 4.2: Summary of synthetic WAFO sequence wave field video attributes for code verification

FFT image (Figure 4.7(b)) is used to determine the maximum energy direction through the slice L. The spatial coordinates of this slice, L, are used throughout the remaining analysis, superimposed on all subsequent frames.



(a)



Figure 4.7: (a) First frame of the *synthetic WAFO sequence* wave field; (b) Direction of maximum energy superimposed on 2D FFT of first scene of the *synthetic WAFO sequene* wave field

To solve for scale, the same slice, L, from the 2-D FFT image per frame is compiled throughout time (left to right) to arrive at Figure 4.8(a). The FFT of each row is then computed and compiled, resulting in Figure 4.8(b). Figure 4.8(b) is used in conjunction with Equations 4.8 – 4.10 in order to arrive at an

image scale, through maximizing the fit of a parabola in the final image (as described in Equation 4.8 and shown in Figure 4.9).



Figure 4.8: (a) 2-D FFT slice compiled image from the *synthetic WAFO sequence;* (b) row FFT computed from each row of (a) and compiled for each row, with resulting best-fit parabola



Figure 4.9: Maximizing α in order to solve for scale, β , for *synthetic*

Figures 4.10 and 4.11 illustrate the steps of the algorithm to recover sea state from a sea surface video. To determine the dominant wave frequency, the energy is averaged, by column, per frequency bin, of each row FFT (Figure 4.11). From this averaged row FFT, the dominant wave frequency can be used to calculate the dominant wave period. With those values, the wavelength can be recovered, along with the wave velocity and maximum wave height.

Results are summarized in Table 4.3. The dominant wave period was recovered exactly, and the scale within 1.5% of the generated *WAFO* video scale


Figure 4.10: (a) Spatio-temporal image of *synthetic WAFO sequence;* (b) row FFT computed from each row of spatio-temporal image (a)



Figure 4.11: Averaged FFT (by column) of Figure 4.10b

(synthetic ground truth scale of 6.7 pixels/meter vs recovered scale of 6.6 pixels/meter). These results constitute an important verification result; thus, now, the reporting moves on to the consideration of real experimental data.

| Sequence | λ [m] | $T\left[\mathbf{s}\right]$ | f_t [Hz] | β [pix/m] | c [m/s] | max H |
|-------------------------|---------------|----------------------------|------------|-----------------|---------|-------|
| | [m] | [s] | [Hz] | [pix/m] | [m/s] | [m] |
| synthetic WAFO sequence | 6.22 | 2.00 | 0.50 | 6.6 | 3.12 | 0.62 |

Table 4.3: Summary of results for wave field in *synthetic WAFO sequence*

4.3 Validation experiment

The present section explores the agreement between physical experiments and the verified simulation source code, in order to more completely understand the predictive power of the algorithm developed by Spencer et al.

4.3.1 Description of data

Wave fields were generated within one of the flumes in residence at the DeFrees Hydraulics Laboratory at Cornell University in Ithaca, New York, in July 2017. Due to the dimensions of the flume and specifications of the wave generator, the flume could only be filled to a depth of 12 inches; thus, the wave heights could not exceed 10mm, or a wave frequency of 1.25 Hz, so as to conform to the assumed "deep water" dispersion relation in Equation 4.4. Table 4.4 outlines the nature of the various wave videos collected from the assessed flume wave fields. Each of the image sequences is labeled with a letter (E, F, G, H, M, and N). In addition to adjusting wave height and frequency, camera grazing angle was alternated between 15 and 30 degrees grazing angles (Figure 4.12), and scene lighting adjusted between overhead lighting versus spotlight lighting (Figure 4.13), respectively. For example, Sequence M has a wave amplitude of 10 mm, a wave frequency of 1.25 Hz, a 15 degree camera grazing angle, and is illuminated by overhead (ambient) lighting, rather than spotlight lighting. Images were collected using a GoPro Hero 5 [18] camera mounted over the top of the flume.

In order for the dispersion relationship (Equation 4.4) to hold, the approximation $kd \approx 2$, where *k* is the wavenumber in Table 4.1 and *d* is the depth of the

| Camera grazing ang | grazing angle 15 degree | | | 30 degree | | | |
|---------------------|-------------------------|----------|-----------|-----------|-----------|--|--|
| Wave amplitude [mm] | Hz | Overhead | Spotlight | Overhead | Spotlight | | |
| 5 | 1.25 | Е | F | G | Н | | |
| 10 | 1.25 | М | Ν | | | | |

Table 4.4: Video collection schedule for small-scale image sequences E, F, G, H, M, and N





Figure 4.12: The two grazing angles used to collect video data: on left, 15° grazing angle; on right, 30° grazing angle

water (12 inches), is applied. This value is chosen such that tanh can be approximated by 1 (Equations 4.1-4.2). The resulting kd value from a 1.25 Hz wave in this wave flume is \approx 2.0, which is accepted to satisfy the approximation [81].

The ground truth for scale, β , was measured by calculating distances between image points of known dimensions on the flume. For example, the cross bars on the bottom of the flume are 16" apart, center to center. Using calmwater images, the number of pixels is determined from one edge of the cross bar to the next. This process is done manually by visual inspection for repeated frames and reference points. The calculated mean, μ , and standard deviation, σ ,



Figure 4.13: The two lighting conditions used to collect video data: on left, ambient overhead lighting; on right, spotlight lighting

for a sample size of 11 repeated image measurements, recorded in Table 4.5. For the cases where there is no perspective correction applied, a standard deviation, σ , is not listed. Instead, a range appears to indicate the approximated range of scale present in the image; the real-world coordinates closer to the camera (at the bottom of the camera frame) will have a larger pixels per meter value in the camera coordinate system, and the real-world coordinates farther from the camera (towards the top of the camera frame) will have a smaller pixels per meter value in the camera coordinate system. Additionally, the perspective correction images result in a smaller scale than the uncorrected image due to the transformation applied to the image.

To validate the wave frequencies and wave heights within the flume, a third camera was placed on the side of the flume at water-level, and recorded simultaneously with the other two cameras. The wave period is confirmed as exactly 0.8s for all video sequences considered (sample size 20). The wave amplitude

| Camera grazing angle | Perspective correction | β , μ [pix/m] | β, σ [pix/m] |
|----------------------|------------------------|-------------------------|-------------------------|
| 15^{o} | none | 172-290 | |
| 15^{o} | overhead | 147.6 | 1.4 |
| 30^{o} | none | 128-315 | |
| 30^{o} | overhead | 139.4 | 1.4 |

Table 4.5: Summary of ground truth scale, β , values

for the nominally 5mm wave amplitude is calculated as an average μ = 5.6 mm (σ = 0.12 mm), and the nominally 10 mm wave amplitude is calculated as an average μ = 10.9 mm (σ = 0.58 mm).

Theoretical limitations

Videos are recorded at 30 or 60 frames per second (fps), well above the minimum Nyquist-Shannon frequency of 2.5Hz required to recover a 1.25 Hz wave. The wavelength of interest is 1m, and all considered video sequences meet the subsequent required scale, $\beta > 2$. The standard sequence length is set to 20 cycles, or 480 frames, in accordance with Spencer et al.'s suggestion [70].

4.3.2 Experimental results

The present section outlines the results from application of Spencer et al.'s wave imaging algorithm to the present flume data.

Raw image results

As a first pass, the algorithm is applied on the image sequence, E, without any post-processing (i.e., image enhancement) for the image frame. Three sequences

are analyzed, labeled *E1*, *E2*, and *E3*. The three sequences are recordings of the same wave field condition obtained using the same camera settings; thus, these are separate instances of the same image sequence. Figure 4.15 shows a demonstration of the single camera image sequence method on an laboratory-generated wave field video scene. The first frame of the sequence is shown in Figure 4.14(a). The resulting 2-D FFT image is shown in Figure 4.14(b). The spatial coordinates of slice, L, is the direction of maximum energy (superimposed on Figure 4.14(b)).





Figure 4.14: (a) First frame of *E1* video; (b) Direction of maximum energy superimposed on 2D FFT of first scene of *E1* video

As described in the synthetic video code verification section above, the same slice, L, from the 2-D FFT image is captured for every in the video sequence, and subsequently compiled (left to right) into an image (Figure 4.15(a)). Applying a FFT for each row in Figure 4.15(a) results in the generation of Figure 4.15(b). Figure 4.15(b) is the final image employed to determine scale (i.e., using Equation 4.9).

Figures 4.16 and 4.17 illustrate the steps of the algorithm to recover sea state



Figure 4.15: (a) *Sequence E1* accumulated 2-D FFT slice composition; (b) row FFT computed from (a) and compiled for each row, with resulting best-fit parabola for *Sequence E1*. Both images have dimensions of *L* rows and *N* columns

from a flume wave field video. To determine the dominant wave frequency, the energy is averaged, by column, per frequency bin, of each row FFT within Figure 4.15(b) to produce Figure 4.17). From this averaged row FFT, the dominant wave frequency can be used to calculate the dominant wave period. The wavelength can be subsequently recovered, along with the wave velocity, and maximum wave height.



Figure 4.16: (a) *Sequence E1* spatio-temporal slice composition image at 30 fps; (b) row FFT computed from (a), computed and compiled for all rows. Both images have dimensions of *L* rows and *N* columns

Results are summarized in Table 4.6. The dominant wave period was recov-



Figure 4.17: Averaged row FFT of Figure 4.16b

ered exactly, but the scale showed fluctuation between the three videos. The ground truth scale for the raw camera perspective is between 172-290 for this video sequence; the results range from 86.1 to 208.2 for the analyzed sequences E1 – E3. The term "raw" in Table 4.6, refers to the fact that no image post-processing enhancement has been applied. The average value listed in the same table is used henceforth as a comparison value for the subsequent wave field analysis.

| Sequence | λ | T | f_t | eta | С | max H |
|-----------------|-----------|------|-------|---------|-------|-------|
| _ | [m] | [s] | [Hz] | [pix/m] | [m/s] | [m] |
| E1, raw | 1.00 | 0.80 | 1.25 | 208.2 | 1.25 | 0.10 |
| E2, raw | 1.00 | 0.80 | 1.25 | 86.1 | 1.25 | 0.10 |
| E3, raw | 1.00 | 0.80 | 1.25 | 111.3 | 1.25 | 0.10 |
| E, raw, average | 1.00 | 0.80 | 1.25 | 135.2 | 1.25 | 0.10 |

Table 4.6: Summary of results for E1–E3, raw. The average value for E, raw, is used henceforth as the baseline comparison value for this wave field

4.3.3 Experimental parameters

Three experimental design parameters were adjusted at the time of video collection: sampling frequency (fps, or frame rate, *v*), camera grazing angle, and lighting condition.

As shown in Figure 4.13, wave field videos were taken under two different lighting conditions. The results for the "spotlight lighting" condition, comprised of Sequences F and H, are presented in Table 4.7. Sequence F is a spotlight sequence recorded at a 15° grazing angle, while Sequence H is recorded at a 30° grazing angle.

As shown in Figure 4.12, wave field videos were also recorded at two different grazing angle conditions. The results for the videos taken from 30° grazing angle (Sequences G and H) are presented in Table 4.7 under the "30° grazing angle" condition. Sequence G is a video sequence recorded at a 30° grazing angle with ambient, overhead lighting, while Sequence H is recorded at a 30° grazing angle using a spotlight for lighting.

Two frame rates are also considered: 30 frames per second and 60 frames per second. The standard, baseline run from Section 4.3.2 is collected at 30 fps. The 60 fps sequences are presented in Table 4.7 under the "60 fps" condition.

Notably, the recovered sea state characteristics for λ , T, f_t , c, and max H are all consistent with the baseline (i.e., data set E, "raw") and ground truth (GT) values.

4.3.4 Image processing parameters

Five image processing parameters were considered as part of a post-processing phase for the image enhancement of the recorded video sequences: perspective correction, histogram balancing, linear contrasting, and image smoothing.

| Condition | Sequence | λ | T | f_t | eta | c | max H |
|------------------|------------|-----------|------|-------|---------|-------|-------|
| | _ | [m] | [s] | [Hz] | [pix/m] | [m/s] | [m] |
| | F1 | 1.00 | 0.80 | 1.25 | 430.3 | 1.25 | 0.10 |
| | F2 | 1.00 | 0.80 | 1.25 | 430.3 | 1.25 | 0.10 |
| Spotlight | F3 | 1.00 | 0.80 | 1.25 | 461.1 | 1.25 | 0.10 |
| lighting | H1 | 1.00 | 0.80 | 1.25 | 99.0 | 1.25 | 0.10 |
| | H2 | 1.00 | 0.80 | 1.25 | 520.6 | 1.25 | 0.10 |
| | H3 | 1.00 | 0.80 | 1.25 | 581.5 | 1.25 | 0.10 |
| | G1 | 1.00 | 0.80 | 1.25 | 75.9 | 1.25 | 0.10 |
| | G2 | 1.00 | 0.80 | 1.25 | 74.2 | 1.25 | 0.10 |
| 30^{o} grazing | G3 | 1.00 | 0.80 | 1.25 | 50.4 | 1.25 | 0.10 |
| angle | H1 | 1.00 | 0.80 | 1.25 | 99.0 | 1.25 | 0.10 |
| | H2 | 1.00 | 0.80 | 1.25 | 520.6 | 1.25 | 0.10 |
| | H3 | 1.00 | 0.80 | 1.25 | 581.5 | 1.25 | 0.10 |
| | E1 | 1.00 | 0.80 | 1.25 | 75.1 | 1.25 | 0.10 |
| 60 fps | E2 | 1.00 | 0.80 | 1.25 | 18.8 | 1.25 | 0.10 |
| | E3 | 1.00 | 0.80 | 1.25 | 55.2 | 1.25 | 0.10 |
| Raw | E, average | 1.00 | 0.80 | 1.25 | 135.2 | 1.25 | 0.10 |
| Ground truth | Е | 1.00 | 0.80 | 1.25 | 173-289 | 1.25 | 0.10 |

Table 4.7: Summary of results for varying experimental design parameters

Perspective introduces potential complications for the image analysis, since multiple scales in the image could be present if the perspective is not corrected. The current study considers a a grazing angle of 15°. To correct for this, a geometric transformation can be applied to the image to recover a geometry representative of an overhead perspective. This can be applied as an affine (shear) transformation, projective transformation, and/or a Euclidean (rotational) transformation. Correctly applied, this transformation should create images that have a uniform scale throughout the image [1]. The perspective-corrected images are labeled as the "Overhead" condition in Table 4.8.

Additionally, two contrast enhancing image processing options are implemented in this study. Both of these affect the intensity of pixels in the resulting post-processed video images. Generally, contrast enhancing can improve image feature identification by optimizing, or modifying, the image grey scale intensity to better distribute intensity values. The first contrast enhancing option considered is *histogram equalization* ("Histeq" in Table 4.8). This correction, applied in MATLAB using function *histeq*, transforms the intensity of a given image such that the distribution of pixels in each intensity bin are evenly distributed. This is not a linear intensity transformation. The second contrast enhancing option considered is a linear intensity transformation that saturates the top 1% and bottom 1% of all pixel values, thereby increasing the contrast of the output image ("linear contrast" in Table 4.8). This is applied in MATLAB through the function *imadjust*.

Image smoothing can help to alleviate any unnecessary image noise, especially at higher frequencies. The specified video frames subjected to smoothing have been filtered with a 2-D Gaussian kernel with a standard deviation of 0.5 pixels, chosen to provide desired filtering properties while maintaining image features. The results are included in Table 4.8, under the "Smoothing" condition.

4.3.5 Algorithm input parameters

Three input parameters were adjusted in the implementation of the algorithm: the direction of maximum energy, image window size (pixels by pixels), and length of sequence (N).

| Condition | Sequence | λ | T | f_t | eta | c | max H |
|-----------------|------------|-----------|------|-------|---------|-------|-------|
| | - | [m] | [s] | [Hz] | [pix/m] | [m/s] | [m] |
| | E1 | 1.00 | 0.80 | 1.25 | 23.1 | 1.25 | 0.10 |
| Overhead | E2 | 1.00 | 0.80 | 1.25 | 82.3 | 1.25 | 0.10 |
| | E3 | 1.00 | 0.80 | 1.25 | 28.3 | 1.25 | 0.10 |
| | E1 | 1.00 | 0.80 | 1.25 | 169.9 | 1.25 | 0.10 |
| Histeq | E2 | 1.00 | 0.80 | 1.25 | 169.9 | 1.25 | 0.10 |
| | E3 | 1.00 | 0.80 | 1.25 | 54.7 | 1.25 | 0.10 |
| | E1 | 1.00 | 0.80 | 1.25 | 88.8 | 1.25 | 0.10 |
| Overhead histeq | E2 | 1.00 | 0.80 | 1.25 | 62.0 | 1.25 | 0.10 |
| | E3 | 1.00 | 0.80 | 1.25 | 38.1 | 1.25 | 0.10 |
| | E1 | 1.00 | 0.80 | 1.25 | 169.9 | 1.25 | 0.10 |
| Linear contrast | E2 | 1.00 | 0.80 | 1.25 | 195.6 | 1.25 | 0.10 |
| | E3 | 1.00 | 0.80 | 1.25 | 87.2 | 1.25 | 0.10 |
| | E1 | 1.00 | 0.80 | 1.25 | 107.6 | 1.25 | 0.10 |
| Smoothing | E2 | 1.00 | 0.80 | 1.25 | 75.1 | 1.25 | 0.10 |
| | E3 | 1.00 | 0.80 | 1.25 | 23.1 | 1.25 | 0.10 |
| Raw | E, average | 1.00 | 0.80 | 1.25 | 135.2 | 1.25 | 0.10 |
| Raw GT | E | 1.00 | 0.80 | 1.25 | 173-289 | 1.25 | 0.10 |
| Overhead GT | E | 1.00 | 0.80 | 1.25 | 147.6 | 1.25 | 0.10 |

Table 4.8: Summary of results for video perspective correction, with and without image enhancement

Direction of maximum energy

In order to consider the sensitivity of results to the selection of the *maximum energy direction* (DME) from which the slice *L* is taken, Table 4.9 summarizes the results of considering eight different maximum energy direction slices considered as perturbations from the calculated DME, ranging from from -1.78° to $+1.78^{\circ}$ off of DME.

Window size

Two additional window sizes are considered, in additional to the original 256 pixel by 256 pixel window: a 64 pixel by 64 pixel window and a 128 by 128 pixel

| Condition | Sequence | λ | T | f_t | eta | c | $\max H$ |
|-------------|------------|-----------|------|-------|---------|-------|----------|
| | | [m] | [s] | [Hz] | [pix/m] | [m/s] | [m] |
| -1.78^{o} | E1 | 1.00 | 0.80 | 1.25 | 107.6 | 1.25 | 0.10 |
| -1.33^{o} | E1 | 1.00 | 0.80 | 1.25 | 208.2 | 1.25 | 0.10 |
| -0.88^{o} | E1 | 1.00 | 0.80 | 1.25 | 38.9 | 1.25 | 0.10 |
| -0.44^{o} | E1 | 1.00 | 0.80 | 1.25 | 208.2 | 1.25 | 0.10 |
| DME | E1 | 1.00 | 0.80 | 1.25 | 208.2 | 1.25 | 0.10 |
| $+0.44^{o}$ | E1 | 1.00 | 0.80 | 1.25 | 222.6 | 1.25 | 0.10 |
| $+0.88^{o}$ | E1 | 1.00 | 0.80 | 1.25 | 146.7 | 1.25 | 0.10 |
| $+1.33^{o}$ | E1 | 1.00 | 0.80 | 1.25 | 230.5 | 1.25 | 0.10 |
| $+1.78^{o}$ | E1 | 1.00 | 0.80 | 1.25 | 215.2 | 1.25 | 0.10 |
| Raw | E, average | 1.00 | 0.80 | 1.25 | 135.2 | 1.25 | 0.10 |
| Raw GT | E | 1.00 | 0.80 | 1.25 | 173-289 | 1.25 | 0.10 |

Table 4.9: Summary of sensitivity results for video E1 at incremental deviations from DME for slice L

window. The method does not appear to have large sensitivities between the original window size and the medium window size, but the small window size does not appear stable (ranging from 133.4 to 640.5 in scale).

| Condition | Sequence | λ | T | f_t | eta | c | max H |
|-----------|------------|-----------|------|-------|---------|-------|-------|
| | | [m] | [s] | [Hz] | [pix/m] | [m/s] | [m] |
| | E1 | 1.00 | 0.80 | 1.25 | 133.4 | 1.25 | 0.10 |
| 64 x 64 | E2 | 1.00 | 0.80 | 1.25 | 640.5 | 1.25 | 0.10 |
| | E3 | 1.00 | 0.80 | 1.25 | 133.4 | 1.25 | 0.10 |
| | E1 | 1.00 | 0.80 | 1.25 | 106.7 | 1.25 | 0.10 |
| 128 x 128 | E2 | 1.00 | 0.80 | 1.25 | 200.2 | 1.25 | 0.10 |
| | E3 | 1.00 | 0.80 | 1.25 | 133.4 | 1.25 | 0.10 |
| Raw | E, average | 1.00 | 0.80 | 1.25 | 135.2 | 1.25 | 0.10 |
| Raw GT | E | 1.00 | 0.80 | 1.25 | 173-289 | 1.25 | 0.10 |

Table 4.10: Summary of results for video sequence E, varying window size

Length of sequence

The Spencer et al. suggest that the wave imaging algorithm requires at least 20 cycles of the wave period captured in a video sequence for reliable accuracy

[70]. The current subsection explores the veracity of the suggestion by using a single cycle, two cycles, five cycles, and ten cycles as comparison values (Table 4.11).

| Condition | Sequence | λ | T | f_t | β | c | max H |
|-----------|------------|-----------|------|-------|---------|-------|-------|
| | _ | [m] | [s] | [Hz] | [pix/m] | [m/s] | [m] |
| 1 | E1 | 1.00 | 0.80 | 1.25 | 2347.2 | 1.25 | 0.10 |
| | E1 | 1.00 | 0.80 | 1.25 | 179.3 | 1.25 | 0.10 |
| 2 | E2 | 1.00 | 0.80 | 1.25 | 134.5 | 1.25 | 0.10 |
| | E3 | 1.00 | 0.80 | 1.25 | 269.0 | 1.25 | 0.10 |
| | E1 | 1.00 | 0.80 | 1.25 | 163.9 | 1.25 | 0.10 |
| 5 | E2 | 1.00 | 0.80 | 1.25 | 169.3 | 1.25 | 0.10 |
| | E3 | 1.00 | 0.80 | 1.25 | 161.4 | 1.25 | 0.10 |
| | E1 | 1.00 | 0.80 | 1.25 | 198.6 | 1.25 | 0.10 |
| 10 | E2 | 1.00 | 0.80 | 1.25 | 112.3 | 1.25 | 0.10 |
| | E3 | 1.00 | 0.80 | 1.25 | 143.4 | 1.25 | 0.10 |
| Raw | E, average | 1.00 | 0.80 | 1.25 | 135.2 | 1.25 | 0.10 |
| Raw GT | E | 1.00 | 0.80 | 1.25 | 173-289 | 1.25 | 0.10 |

Table 4.11: Summary of results for video sequence E, ranging from 1 observed cycle to 10 observed cycles. The raw video sequence observed 20 cycles

4.3.6 Y-aware parameter selection

The previous set of prescribed image capture parameters (e.g., L, N, v) were chosen based on recommendations from Spencer [70], combined with physical limitations of the available experimental facility. Using these prescribed parameters, it was possible to accurately extract the frequency content from the wave field video, regardless of camera graze angle or any of the foregoing image processing methods employed. Though the scale was accurately recovered for the synthetic data set, the scale recovery for the validation experiment (carried out in the flume) has, thus far, shown less consistency.

Though Spencer et al. recommends parameters for the number of cycles captured, window sizes considered for analysis, and spatial or temporal frequencies required for reliable analysis, the interplay between these factors in producing a stable β solution is not carefully examined [70]. For the case of non-synthetic data – in this case, real world data – the parabola of interest (Equation 4.10) is more susceptible to noise, and the value for α can be difficult to extract from the FFT of the spatial frequency temporal image (as depicted in Figure 4.2).

Beyond the empirically observed sensitivity of β to measurement noise, the value of β is also proportional to the inverse value of α (Equation 4.10). This means that, as α approaches zero, the value of β approaches infinity (i.e., the formula for β is *ill-conditioned* on α). Taking a closer look at Equation 4.10, the coefficient of the $1/\alpha$ term is dependent on parametric values chosen by the analyst. For example, frame rate, v, the length of the sequence, N, and the length of slice, L, are all within control of the analyst implementing the algorithm, and special attention must be given to the relationship between these values. For ease, a new coefficient, Y, is defined as the following:

$$\beta = \frac{2\pi L v^2}{g N^2} \frac{1}{\alpha} = \frac{Y}{\alpha}$$
(4.11)

With the original parameters, L = 258, N = 480, v = 30, the coefficient of the $1/\alpha$ term, Y, is 0.643 (including the constant terms). Given that the ground truth for E, raw, is approximately $\beta = 230$, the accompanying α term must equal 0.0028, and any small perturbation in α at that relatively small magnitude can result in a large difference to β . A longer video sequence can introduce large computational errors when solving for the α term, since the length of the sequence N is squared and in the denominator. Alternatively, more frames per

second, or a larger pixel window, can alleviate the need for α to be very small. This can help explain why the method appears to be more accurate for β values that are relatively small.

Due to the sensitivities and effects of the parameters L, N, and v on the value of α , the authors propose a modified selection process of these parameters. In addition to prescribing the number of wave cycles captured per video duration, or wavelengths present per slice L, the authors suggest that the Y coefficient in Equation 4.11 should be selected in order to satisfy a minimum value. This minimum value should be chosen such that it allows α to remain within a region where there is still reasonable variation, but without introducing unwanted illconditioning on β . For example, the difference in β between 1/1000 (where α = 1000) and 1/1001 (where α = 1001) is insignificant and would require a search over a large range of α values, whereas the resulting change between 1/0.00001 and 1/0.000011 could have a huge effect on β . Thus, the current study recommends the parameters L, N, and v be modified to L = 258, N = 120, v = 30. The Y value for the raw video is now equal to 10.3 (for perspective corrected videos, the coefficient Y is now 7.24). The α value required to recover a scale of 230 is now 0.045, rather than 0.0028. This requires an estimate on the order of magnitude that is expected in the scale of the video, but in general, the larger the coefficient value, the better for fitting an appropriate α value.

To explore the proposed, new parameter selection, based on the Y coefficient, a majority of the previous set of analyses is repeated with the new parameters applied. Previous experiments related to parameters L, N, and v are not repeated, as these are now fixed by virtue of a prescribed Y coefficient. Specifically, the following set includes results for conditions of lighting, camera angle,

perspective correction, contrast through histogram equalizing, contrast through linear adjustment, image smoothing, and direction of maximum energy. Since the results for all sea state characteristics are recovered successfully, only α and β are presented in the following results.

Y-aware Sequence E

The new "raw" values for the same Sequence E, now analyzed under the *Y*-aware selection parameters, show an improvement in consistency of scaling results, as well as improved accuracy when compared to the ground truth. Table 4.12 summarizes the results from this image sequence, E, without any image post-processing and using the *Y*-aware parameter selections. The α value is the optimized α value from Equation 4.9, and the β value is calculated from Equation 4.10, furnished with the indicated α value.

| Sequence | α | β [pix/m] |
|-------------------------|----------|-----------------|
| E1 Raw, Y-aware | 0.0628 | 164.5 |
| E2 Raw, <i>Y</i> -aware | 0.0613 | 168.5 |
| E3 Raw, Y-aware | 0.0643 | 160.6 |
| Baseline average | 0.0628 | 164.5 |

Table 4.12: Summary of results for video E1, E2, and E3, with *Y*-aware parameter selection

Experimental parameters

Experimental design parameters of lighting and camera grazing angle are also now re-considered. Overall, the *Y*-aware sequences show an improvement in the stability of the scale solutions amongst available videos.

| Condition | Sequence | α | β [pix/m] |
|------------------|------------|----------|-----------------|
| | F1 | 0.0452 | 228.5 |
| | F2 | 0.0999 | 103.4 |
| Spotlight | F3 | 0.0733 | 140.9 |
| lighting, | H1 | 0.0412 | 250.7 |
| <i>Y</i> -aware | H2 | 0.0677 | 152.6 |
| | H3 | 0.1182 | 87.4 |
| | G1 | 0.1759 | 58.7 |
| | G2 | 0.0418 | 247.1 |
| 30^{o} grazing | G3 | 0.1252 | 82.5 |
| angle, | H1 | 0.0412 | 250.7 |
| <i>Y</i> -aware | H2 | 0.0677 | 152.6 |
| | H3 | 0.1182 | 87.4 |
| Raw, Y-aware | E, average | 0.0628 | 164.5 |

Table 4.13: Summary of results for Experimental parameters, *Y*-aware

4.3.7 Image processing parameters, *Y*-aware

The imaging processing selections considered in this subsection include perspective correction, contrast through histogram equalizing, contrast through linear adjustment, and image smoothing. Results from consideration of these image processing selections are summarized in Table 4.14.

The perspective corrected sequences ("Overhead" condition in Table 4.14) show an improvement in the accuracy of the scale solutions. Previously the β value ranged from 23 to 82, and with the *Y*-aware parameters, range from 54 to 107. Taking a closer look at the 54 value, the α search is shown in Figure 4.18. All three sequences have a peak around 0.07, and all three sequences have a few local maxima, but in the third (bottom) sequence, the global maximum is inconsistent with the global maximum value of the other two sequences. This results in a different value of α chosen to use in the calculation of β .

It is subsequently pointed out that in some instances, it may be appropriate to choose the first, local maximum within the parabola energy versus α plot, especially in situations where the global and local maxima are somewhat equivalent in magnitude (such as in the case of Figure 4.18).



Figure 4.18: α search for Sequence E1 (top), E2 (middle), and E3 (bottom)

| Condition | Sequence | α | β [pix/m] |
|-----------------|------------|----------|-----------------|
| | E1 | 0.0693 | 105.1 |
| Overhead | E2 | 0.0681 | 107.0 |
| | E3 | 0.1345 | 54.2 |
| | E1 | 0.0504 | 204.9 |
| Histeq | E2 | 0.0594 | 173.9 |
| | E3 | 0.0603 | 171.3 |
| | E1 | 0.0859 | 84.8 |
| Overhead histeq | E2 | 0.0703 | 103.6 |
| | E3 | 0.0663 | 109.9 |
| | E1 | 0.0504 | 204.9 |
| Linear contrast | E2 | 0.0525 | 196.7 |
| | E3 | 0.0653 | 158.2 |
| | E1 | 0.0487 | 212.1 |
| Smoothing | E2 | 0.0664 | 155.5 |
| _ | E3 | 0.0656 | 157.4 |
| Raw | E, average | 0.0628 | 164.5 |

Table 4.14: Summary of results for video perspective correction, alongwith image enhancement through post-processing, for Y-
aware parameter selection

Direction of maximum energy

The directional study now shows less sensitivity to the direction of maximum energy (DME) than before, though the solution still shows a range of results. Directions from -1.78° to $+1.78^{\circ}$ relative to the actual direction of maximum energy are varied under the *Y*-aware parameter selection analyses. Results are summarized in Table 4.15.

| Condition | Sequence | α | β [pix/m] |
|-------------|------------|----------|-----------------|
| -1.78^{o} | E1 | 0.0521 | 198.2 |
| -1.33^{o} | E1 | 0.0878 | 117.6 |
| -0.88^{o} | E1 | 0.0615 | 167.9 |
| -0.44^{o} | E1 | 0.0499 | 207.0 |
| DME | E1 | 0.0878 | 117.6 |
| $+0.44^{o}$ | E1 | 0.0350 | 295.1 |
| $+0.88^{o}$ | E1 | 0.0502 | 205.7 |
| $+1.33^{o}$ | E1 | 0.0502 | 205.7 |
| $+1.78^{o}$ | E1 | 0.0487 | 212.1 |
| Raw | E, average | 0.0628 | 164.5 |

Table 4.15: Summary of results for video E1 at varying directions for slice *L*, labeled as difference in degrees from true direction of maximum energy (DME), for the *Y*-aware parameter selection

4.3.8 Sequence M

The *Y*-aware selection process is demonstrated again on a different sequence (M, rather than E), with different parameters selected for N, v, and L, but with the same resulting *Y* coefficient. Sequence M also contains waves of twice the previous wave height, nominally 20mm rather than 10mm. A subset of the previous validation, considering only the image processing parameters and implementing the *Y*-aware selection criteria at the expense of the number cycles

captured in the video, is presented in this section. This set again features waves with a frequency of 1.25Hz and the same sea state characteristics and image scales as before. However, the video is recorded at 60 fps, and again only five cycles are captured (total of 240 frames). Again, since sea state characteristics are accurately recovered, only values for α and β are presented in the following tables.

M baseline

Results for the *M* baseline are summarized in Table 4.16. The *Y*-aware criteria for parameter selection did not perform as accurately or consistently as it had on the previous dataset (i.e., data set *E*), which has the same *Y* value, but differed in frames per second and in number of frames. Taking a closer look at the α values, the α search is shown in Figure 4.19. All three sequences have a peak around 0.09, and all three sequences have a few local maxima, but in the third (bottom) sequence, the largest peak is not the first peak. This results in a different value of α chosen to use in the calculation of β . This further supports the previous finding: on occasions where the magnitude of the local and global maxima are close, it may be desirable to select the first such local maximum with which to compute β .

| Sequence | α | β [pix/m] |
|----------|----------|-----------------|
| M1 raw | 0.0916 | 112.7 |
| M2 raw | 0.1941 | 53.2 |
| M3 raw | 0.1025 | 100.8 |

Table 4.16: Summary of results for Sequence M, *Y*-aware



Figure 4.19: Alpha search for Sequence M1 (top), M2 (middle), and M3 (bottom)

4.3.9 Image processing parameters, Sequence M

The following results for Sequence M (Table 4.17) are provided for completeness, but all of them exhibit a similar trend to the baseline analysis. Though the general shapes are consistent for finding the desired scale value, there is substantial noise in the search for the best-fit α from which to calculate the scale, β .

| Condition | Sequence | α | β [pix/m] |
|-----------------|----------|----------|-----------------|
| | M1 | 0.0833 | 73.0 |
| Overhead | M2 | 0.1795 | 33.9 |
| | M3 | 0.0812 | 74.9 |
| | M1 | 0.0802 | 128.8 |
| Histeq | M2 | 0.0784 | 131.7 |
| - | M3 | 0.0571 | 180.9 |
| | M1 | 0.1165 | 52.2 |
| Overhead histeq | M2 | 0.1130 | 53.8 |
| _ | M3 | 0.0190 | 320.2 |
| | M1 | 0.0831 | 124.3 |
| Linear contrast | M2 | 0.0792 | 130.4 |
| | M3 | 0.1819 | 56.8 |
| | M1 | 0.3173 | 32.5 |
| Image smoothing | M2 | 0.1975 | 52.3 |
| - 0 | M3 | 0.0985 | 104.9 |

Table 4.17: Summary of results image processing parameters, *Y*-aware

4.4 Discussion

Despite altering camera positioning, lighting conditions, and image processing implemented, all the videos reliably recovered the accurate frequency content of the water videos. This meant that the associated dominant wave period, wavelength, maximum wave height, and wave velocity could also be recovered from the wave video.

The original algorithm, as prescribed by Spencer et al. [70] calls attention to the choice of parameters of video length and window size, as well as sampling frequency and spatial frequency. However, Spencer et al. [70] did not discuss the implications of balancing choices between the number of frames, N, frames per second v, and slice length L (related to window size).

The present work suggests that these parameters should be chosen specifically, with attention to the *Y* coefficient, to ensure that α can remain above an identified value, so as to ensure accuracy and consistency in estimation of the spatial scale factor, β (i.e., to avoid numerical ill-conditional). While it is tempting to presume that more data (i.e., longer video sequences) will lead to higher accuracy, the ensuing ill-conditions may increase the computational demand in the calculation of β .

The experimental dataset itself has shortcomings that require future attention. The first is that a flume was used to generate water waves. As shown in Figures 4.12 and 4.13, the bottom of the flume is visible in the images of the water (i.e., in the open sea). This likely contributes to the horizontal lines in the 2D-FFT of the still frame image, used to determine the direction of maximum energy. Additionally, the water waves are not representative of how they would appear in the conditions of interest. This may be improved through the use of some sort of dye or suspension in the water, such that the bottom of the flume would not be visible. Unfortunately, due to the specifications of the flume and wave maker, the maximum water depth is prescribed to be no greater than 12 inches. The wavelengths and wave heights are inherently limited by the water depth, if the dispersion relationship (Equation 4.3) is to remain valid. Thus, a full demonstration of waves of a real-world size could not be investigated within the scope of this study. It also meant that the study could only assess a single wave frequency, as the requirements were near the very edge of the experimental facility's capabilities.

In the FFT of the spatial-temporal image, periodic vertical lines are visible (see Figure 4.15). An additional step of processing had been applied to explore whether these lines impact the results of the study. This noise was removed by taking the 2D-FFT of the image (Figure 4.15), masking the portion of the image that corresponded to the lines, and then taking the inverse 2D-FFT of the masked image (Figure 4.20). This alternate pre-processing method was then re-run with four random trials above. The results did not change significantly, with β changing less than 10% for each individual sequence per condition. For comparison, this variation is less than the variation within sequences. Thus, this was not pursued further in the current study.



Figure 4.20: The FFT of the spatial-temporal image (Figure 4.15), filtered to minimize vertical lines via masking the 2-D FFT

4.5 Conclusion

The originally proposed algorithm of Spencer et al. [70] seeks to recover both sea-state and image scale in pixels per meter from a single-perspective, uncalibrated video scene of water. Spencer et al. [70] verify the algorithm extensively through the use of synthetic wave videos. In order to further understand how this method translates from synthetic data to the real-world, the present work seeks to validate the algorithm through a suite of experimentally collected video sequences. Prior to presenting the validation experiment, an overview of the original method, and a demonstration of the code verification, is provided to build confidence in the implementation of the intended algorithm.

The application of an uncalibrated wave video to identify scale and sea state is promising though somewhat limiting due to visibility in low-light (night), or obstructed visibility (such as fog). The logistics of directionality would also have to be considered in order to obtain a panoramic view. However, it provides an economical alternative to monitoring of sea states. Additionally, since it requires only a single camera, the typical concerns of calibrating for camera extrinsics, relative to one another, can be entirely eliminated. In this way, the method promises an aspect of enhanced robustness as opposed to some stereo camera methods requiring careful hardware calibration and image registration.

APPENDIX A Chapter 1 of Appendix

A.1 Markov chain Monte Carlo

The following details the theory and specifics of a Markov chain Monte Carlo inversion applied in this study, specifically using the Metropolis-Hastings algorithm.

The inversion seeks to "build" a chain of values to provide an estimate for $p(k^{roll}|T_{GT})$ in the form of a first order Markov chain, which is written as $f(k_n^{roll}|k_{n-1}^{roll})$. Any current value at n is conditional on only the past value, n - 1. The Metropolis-Hastings algorithm is applied to accept or reject proposed moves from a random walk through the parameter space (*e.g.*, k^{roll}) in order to build this chain for k_n^{roll} , where n has values from 1 to a user-prescribed N. Hastings outlines this approach in detail in [23]. The algorithm requires that the Markov chain be stationary, symmetric, and that each of the values in the sequence are sampled from the same distribution, $p(k^{roll})$, for detailed balance.

The Metropolis-Hastings algorithm is implemented to compute, at any step n, a trial value of the roll gyradius, called k_n^{roll} . From k_n^{roll} in a forward model, T_n is attained. T_n is analogous to the value earlier denoted as $SMP(k^{roll})$, now specifically $T_n = SMP(k_n^{roll})$. The trial T_n can then be compared with the previous T_{n-1} , and assessment of whether k_n^{roll} is a value that should be included in the chain is made. After this comparison, k_n^{roll} is then either included in the chain or over-written by k_{n-1}^{roll} , such that $k_n^{roll} = k_{n-1}^{roll}$. A small "step" is taken away from value k_n^{roll} to a new value, k_{n+1}^{roll} , and repeat the process. This is con-

tinued this for a user-prescribed number of trials until we have a chain of values that satisfies the condition that, for a Markov chain of length O:

$$f(k_n^{roll}|k_{n-1}^{roll}, k_{n-2}^{roll}, ..., k_{n-O}^{roll}) = f(k_n^{roll}|k_{n-1}^{roll}, k_{n-2}^{roll}, ..., k_{n-O}^{roll}, k_{n-O-1}^{roll}, k_{n-O-2}^{roll}, ...)$$

such that obtaining the current value at *n* is dependent only on the last O values, and not values beyond. In this case, the chain only depends on step n-1.



Figure A.1: Inverse model – building the Markov chain

To actually compare the trial value at step n with the value at n - 1, the ratio of likelihood functions must be computed. The ratio of the likelihood functions for this purpose takes the form:

$$r(k_n^{roll}|k_{n-1}^{roll}) = \frac{p(k_n^{roll}|T_{GT})}{p(k_{n-1}^{roll}|T_{GT})}$$
(A.1)

as described below.

An immediate way to compute either $p(k_n|T_{GT})$ or $p(k_{n-1}|T_{GT})$ is not available. However, from Bayes' Theorem (recall Equation 2.3), the following form results:

$$r(k_n^{roll}|k_{n-1}^{roll}) = \frac{p(T_{GT}|k_n^{roll})p_{prior}(k^{roll})/p(T_{GT})}{p(T_{GT}|k_{n-1}^{roll})p_{prior}(k^{roll})/p(T_{GT})}$$
(A.2)

Terms that appear in both the numerator and denominator (symmetric and stationary PDF $p_{prior}(k^{roll})$ and $p(T_{GT})$) can be eliminated. The following is produced:

$$r(k_n^{roll}|k_{n-1}^{roll}) = \frac{p(T_{GT}|k_n^{roll})}{p(T_{GT}|k_{n-1}^{roll})}$$
(A.3)

and all the resources are available to obtain a value for $r(k_n^{roll}|k_{n-1}^{roll})$, denoted also simply as r. Substituting appropriately from Equation 2.4 produces the result:

$$r(k_n^{roll}|k_{n-1}^{roll}) = \frac{\frac{1}{\sqrt{2\pi\sigma_T^2}} e^{-\frac{1}{2\sigma_T^2}(T_{GT} - T_n)^2}}{\frac{1}{\sqrt{2\pi\sigma_T^2}} e^{-\frac{1}{2\sigma_T^2}(T_{GT} - T_{n-1})^2}}$$
(A.4)

which can reduce further to

$$r(k_n^{roll}|k_{n-1}^{roll}) = e^{-\frac{1}{2\sigma_T^2}[Q_n - Q_{n-1}]}$$
(A.5)

where $Q_n = (T_{GT} - T_n)^2$ and $Q_{n-1} = (T_{GT} - T_{n-1})^2$.

If r is greater than 1, the candidate is accepted at step n. If r is less than 1, then a "biased coin toss" is performed. Producing a uniformly distributed random number between 0 and 1, $U_n[0,1]$, r is compared to this number: if r is greater than U_n , then k_n^{roll} is accepted, even though it is not "more likely" than the previous candidate; if r is less than U_n , then we reject k_n^{roll} .

Now that it is known how to compute and how to utilize the ratio in likelihood functions, the remaining component in the method hinges on generating candidates to be included in the random walk.

As alluded to above, after computing the ratio of the likelihood function and either accepting or rejecting the new value, we move forward by taking a small "step" away from value k_n^{roll} to a new value, k_{n+1}^{roll} .

$$k_{n+1}^{roll} = k_n^{roll} + \alpha * U_n[-1, 1]$$
(A.6)

where $U_n[-1,1]$ is a uniformly distributed random number between -1 to 1.

The next question that has yet to be answered is: how is a value for α obtained? To address this, the inversion will now be specified as two stages: a "burn-in" stage and a "sampling" stage.

During the "burn-in" stage, sample collection is started with reasonable diversity within the support of the prior distribution. At this time, discarded samples are used to tune the α parameter such that reasonable rates of acceptance are achieved while maintaining reasonably small auto correlations within the sampling chain when the actual chain sampling is carried out, after the "burn-in" phase. As one means of achieving these outcomes, we adopt the approach of Link [41].

If $r_n \ge U_n[0, 1]$, while $n < N_{burn}$,

$$\alpha_n = \alpha_{n-1} * \left[a - (a-1)\frac{n}{N_{burn}}\right] \tag{A.7}$$

Else, if $r_n < U_n[0, 1]$, while $n < N_{burn}$,

$$\alpha_n = \alpha_{n-1} \frac{1 - P_{ac}[a - (a-1)\frac{n}{N_{burn}}]}{1 - P_{ac}}$$
(A.8)

The parameters P_{ac} (the target acceptance rate) and a (tuning parameter) are set to equal 0.42 and 1.01, respectively, as recommended by Link when searching over a single parameter. Roll gyradius values from the "burn-in" stage are not stored or included in the final estimation of $p(k^{roll})$. Now with the inverse model background, the description of the computational framework can be developed to implement such a model (A.2). Algorithm 1: MCMC Stochastic Solver

1: procedure STOCH_SOLVE(N_{burn} , N_{samp} , prior_support, σ_T , T_{GT})

2: [alpha_opt, data, $T_{N_{burn}}$] = burnin(N_{burn} , prior, σ_T , T_{GT})

> Returns burned-in data chain, tuned

alpha, and $T_{N_{burn}}$ 3: data = sample(N_{samp} , prior, alpha_opt, σ_T , T_{GT} , $T_{N_{burn}}$,

 N_{burn} , data) \triangleright data sampling

4: return(data)

5: end procedure

Algorithm 2: Burn-in chain and tuning of alpha

| 1: | 1: procedure BURNIN(N_{burn} , prior, σ_T , T_{GT}) | | | | |
|-----|---|---|--|--|--|
| | Initialize prev_resp, linkset, data, and alpha structures | | | | |
| 2: | linkset = mean(prior) | ▷ Mean of uniform prior | | | |
| 3: | alpha = 0.1 | Starting alpha value | | | |
| 4: | Pac = 0.42 | Darget acceptance rate | | | |
| 5: | a = 1.01 | > Target convergence target | | | |
| 6: | for $n = 1:N_{burn}$ do | | | | |
| 7: | [linkset, conv, prev_resp] = Gibbs(prior, alpha, σ_T , T_{GT} , prev_resp, | | | | |
| | linkset) | ▷ Use Gibbs to handle multiple param- | | | |
| 8: | if r_n is accepted then | eters if needed \triangleright conv is equal to 1 if r_n is accepted | | | |
| 9: | 9: alpha= alpha * $a - (a - 1)\frac{i}{N_{burn}}$ | | | | |
| 10: | else | | | | |
| 11: | 1: alpha = alpha / $\frac{1-P_{ac}}{1-P_{cc}\left(a-(a-1)\frac{i}{a-a}\right)}$ | | | | |
| 12: | 2: end if | | | | |
| 13: | end for | | | | |
| 14: | alpha₋opt = alpha | | | | |
| 15: | return(alpha_opt, data, prev_resp) | | | | |
| 16: | 16: end procedure | | | | |

Algorithm 3: Sampling MCMC for storage

- 1: **procedure** SAMPLE(N_{samp} , prior, alpha_opt, σ_T , T_{GT} , prev_resp, N_{burn} , data)
- 2: linkset = data[N_{burn}]
- 3: $alpha = alpha_opt$
- 4: **for** i in range(N_{burn} , (N_{burn} + N_{samp}),1) **do**
- 5: [linkset, conv, prev_resp] = Gibbs(prior, alpha, σ_T , T_{GT} , prev_resp, linkset)
- 6: Append linkset to data
- 7: end for
- 8: return(data)
- 9: end procedure

Algorithm 4: Allow for capability to vary multiple parameters through Gibbs sampling

```
1: procedure GIBBS(k, alpha, \sigma_T, T_{GT}, prev_resp, linkset)
```

- 2: conv = [0] * length(linkset) > zero-out conv
- 3: **for** i = 1:length(linkset) **do**
- 4: param_min_i = params[2*i] ▷ set minimum from prior
- 5: $param_max_i = params[2^{i+1}] > set maximum from prior$
- 6: [link, conv_flag, prev_resp] = MH(param_min_i, param_max_i,

```
alpha, \sigma_T, T_{GT}, prev_resp, linkset, paramnum)
```

- 7: linkset[i] = link
- 8: conv[i] = conv_flag
- 9: end for
- 10: return(linkset, conv, prev_resp)
- 11: end procedure

Algorithm 5: Performs Metropolis-Hastings move on given link in chain

procedure MH(param_min_i, param_max_i, alpha, σ_T, T_{GT}, prev_resp, linkset, paramnum)

> Generate a candidate move

- 2: candidate = linkset[paramnum] + alpha[paramnum] * rand(-1,1)
 > Test to see if candidate is admissible under prior, return if not
- 3: **if** (**then**candidate >param_max_i or candidate <param_min_i):
- 4: link = linkset[paramnum]
- 5: $\operatorname{conv}_{\operatorname{flag}} = 0$
- 6: return link, conv_flag, prev_resp
- 7: end if

Call forward model to compute the response of candidates that are acceptable under the prior

Initialize temporary data structures p and q

9: p = [0]*len(prev_resp)

10:
$$q = [0]$$
*len(prev_resp)

> evaluate likelihood ratio

- 11: for i in range(len(prev_resp)):
- 12: $p[i] = expt[i] prev_resp[i]$
- 13: $q[i] = expt[i] new_resp[i]$
- 14: $Q_prev = sum(p^*p \text{ for } p, p \text{ in } zip(p, p))$

15: $Q = sum(q^*q \text{ for } q,q \text{ in } zip(q,q))$

16: $r = e^{\frac{-0.5}{\sigma_T}Q - Q_prev}$
| 17: | if r >rand(0,1) then ▷ perform biased coin toss |
|-----|---|
| 18: | link = candidate |
| 19: | prev_resp = new_resp |
| 20: | $conv_flag = 1 \triangleright flag r as accepted$ |
| 21: | else |
| 22: | link = linkset[paramnum] |
| 23: | $conv_flag = 0 \triangleright flag r as rejected$ |
| 24: | return(return link, conv_flag, prev_resp) |
| 25: | end if |
| 26: | end procedure |

Algorithm 6: Running the forward model

- 1: **procedure** FORMOD(linkset, paramnum, candidate, prev_resp)
- 2: Initialize the parameter data structure
- 3: Initialize the new response data structure
- 4: Identify and open the 'base' file name containing the input file for the forward model
- 5: Identify and open the current file name containing the input file for the forward model at step n
- 6: Modify the 'base' file such that the k^{roll} input value equals k_n^{roll} and save as the current file name
- 7: Close all files
- 8: Run SMP with the current file
- 9: Save the output
- 10: Extract roll period T_n in seconds and save as new_response
- 11: return(new_response)
- 12: end procedure

APPENDIX B

CHAPTER 2 OF APPENDIX

B.1 Additional tables for icing configuration results

| | Prior Type | | | | | |
|---------------|------------------------------|----------|---------|---------|----------|----------|
| Prior Support | Variable | | (1) U-U | (2) G-G | (3) Un-G | (4) G-Un |
| Un-Iced | k _{roll} [B] | μ | 0.4107 | 0.3934 | 0.6383 | 0.3865 |
| | | σ | 0.0001 | 0.0059 | 0.0083 | 0.0050 |
| | VCG [m] | μ | 6.4737 | 6.7789 | 6.1830 | 6.7922 |
| | | σ | 0.0002 | 0.0113 | 0.0257 | 0.0094 |
| Experiment | <i>k</i> _{roll} [B] | μ | 0.3745 | 0.3764 | 0.5724 | 0.3774 |
| | | σ | 0.0105 | 0.0055 | 0.0278 | 0.0052 |
| | VCG [m] | μ | 6.8132 | 6.8100 | 6.3711 | 6.8082 |
| | | σ | 0.0185 | 0.0100 | 0.0750 | 0.0096 |
| SPII | k _{roll} [B] | μ | _ | _ | _ | _ |
| | | σ | _ | _ | _ | _ |
| | VCG [m] | μ | _ | _ | - | _ |
| | | σ | _ | _ | - | _ |
| Hybrid | k _{roll} [B] | μ | _ | _ | 0.5559 | _ |
| | | σ | - | _ | 0.0232 | — |
| | VCG [m] | μ | _ | _ | 6.4155 | _ |
| | | σ | _ | _ | 0.0609 | — |

Table B.1: Summary of full-scale from model-scale MCMC inversion results for Icing Configuration 2

| | Prior Type | | | | | |
|---------------|------------------------------|----------|---------|---------|----------|----------|
| Prior Support | Variable | | (1) U-U | (2) G-G | (3) Un-G | (4) G-Un |
| Un-Iced | <i>k</i> _{roll} [B] | μ | 0.3927 | 0.3917 | 0.4170 | 0.3913 |
| | | σ | 0.0116 | 0.0051 | 0.0303 | 0.0047 |
| | VCG [m] | μ | 6.0509 | 6.0564 | 5.9153 | 6.0588 |
| | | σ | 0.0605 | 0.0279 | 0.1648 | 0.0257 |
| Experiment | <i>k</i> _{roll} [B] | μ | 0.3909 | 0.3856 | 0.2892 | 0.3897 |
| | | σ | 0.0087 | 0.0055 | 0.0176 | 0.0049 |
| | VCG [m] | μ | 6.0614 | 6.0907 | 6.5104 | 6.0675 |
| | | σ | 0.0439 | 0.0300 | 0.0665 | 0.0265 |
| SPII | <i>k</i> _{roll} [B] | μ | 0.3848 | _ | _ | 0.3811 |
| | | σ | 0.0023 | _ | _ | 0.0018 |
| | VCG [m] | μ | 6.0905 | _ | - | 6.1115 |
| | | σ | 0.0118 | - | - | 0.0126 |
| Hybrid | k_{roll} [B] | μ | 0.3775 | _ | 0.4040 | 0.3817 |
| | | σ | 0.0115 | - | 0.0167 | 0.0050 |
| | VCG [m] | μ | 6.1289 | _ | 5.9911 | 6.1080 |
| | | σ | 0.0583 | _ | 0.0885 | 0.0261 |

Table B.2: Summary of full-scale from model-scale MCMC inversion results for Icing Configuration 3

| | Prior Type | | | | | |
|---------------|------------------------------|----------|---------|---------|----------|----------|
| Prior Support | Variable | | (1) U-U | (2) G-G | (3) Un-G | (4) G-Un |
| Un-Iced | <i>k</i> _{roll} [B] | μ | 0.4106 | 0.3930 | _ | _ |
| | | σ | 0.0002 | 0.0051 | _ | — |
| | VCG [m] | μ | 6.3933 | 6.5887 | - | _ |
| | | σ | 0.0007 | 0.0143 | - | - |
| Experiment | k _{roll} [B] | μ | _ | _ | _ | _ |
| | | σ | _ | _ | _ | _ |
| | VCG [m] | μ | _ | _ | - | _ |
| | | σ | _ | _ | _ | _ |
| SPII | k _{roll} [B] | μ | 0.4905 | _ | _ | _ |
| | | σ | 0.0005 | _ | _ | _ |
| | VCG [m] | μ | 6.6429 | _ | _ | — |
| | | σ | 0.0011 | _ | - | _ |
| Hybrid | k _{roll} [B] | μ | _ | 0.4557 | _ | _ |
| | | σ | — | 0.0046 | — | — |
| | VCG [m] | μ | _ | 4.9573 | - | — |
| | | σ | - | 0.0474 | _ | - |

Table B.3: Summary of full-scale from model-scale MCMC inversion results for Icing Configuration 6

B.2 Additional icing configuration results, by prior forms

| Symbol | Meaning |
|----------------|---|
| Grey rectangle | k_{roll} and VCG ground truth values, range of $+/-2\sigma$ |
| | (based on exerimental full-scale from model-scale values) |
| blue "." | μ for k_{roll} and VCG ground truth values |
| | (based on exerimental full-scale from model-scale values) |
| black "o" | Inverse solution of k_{roll} and VCG for category of interest |
| | (Category of interest listed in the plot title) |
| black "." | Inverse solution of k_{roll} and VCG for all other categories |
| | Table B.4: Legend for Figures 3.10 – 3.13 |



Figure B.1: Icing Configuration 2 Results: By prior forms



Figure B.2: Icing Configuration 3 Results: By prior forms



Figure B.3: Icing Configuration 6 Results: By prior forms

Un-Iced Prior Support Experiment Prior Support 7 ACG [m] 7 7 7 6 [m] • 🖸 Θ 5 5 0.5 0.3 0.5 0.3 0.4 0.6 0.4 0.6 k_{roll} [B] k_{roll} [B] **Hybrid Prior Support SPII Prior Support** 7 ACG [m] 7 6 0 0 Θ 0 5 5 0.4 0.5 0.6 0.5 0.3 0.3 0.4 0.6 k_{roll} [B] k_{roll} [B]

B.3 Additional icing configuration results, by prior support values

Figure B.4: Icing Configuration 2 Results: By prior support



Figure B.5: Icing Configuration 3 Results: By prior support



Figure B.6: Icing Configuration 6 Results: By prior support

B.4 Additional 2-D icing configuration results



Figure B.7: Icing Configuration 0, Gaussian-Gaussian Form, Experiment Prior Support



Figure B.8: Icing Configuration 0, Gaussian-Gaussian Form, Un-Iced Prior Support



Figure B.9: Icing Configuration 0, Uniform-Uniform Form, Experiment Prior Support



Figure B.10: Icing Configuration 0, Uniform-Uniform Form, Un-Iced Prior Support



Figure B.11: Icing Configuration 5, Gaussian-Gaussian Form, Experiment Prior Support



Figure B.12: Icing Configuration 5, Gaussian-Gaussian Form, Un-Iced Prior Support



Figure B.13: Icing Configuration 5, Uniform-Uniform Form, Experiment Prior Support



Figure B.14: Icing Configuration 5, Uniform-Uniform Form, Un-Iced Prior Support

APPENDIX C

CHAPTER 3 OF APPENDIX

C.1 Convergence criterion

In choosing the number of "burn-in" and "stored" samples, ensuring that the MCMC sampling chain has sufficiently converged is critical, given the specific framework and model. A convergence criterion can guide the choice of "burn-in" and "stored" sample values that the framework requires as a user-input. This builds confidence in the results of the generated posterior for a given mass parameter, θ .

To monitor convergence within the MCMC sample population, a general convergence diagnostic method by Brooks and Gelman is implemented (Brooks:1998). This method is favorable because it does not require the assumption of any particular distribution (Gaussian or otherwise), and is easily applied to computational models. The method is as follows: start with M chains of length N. For a given chain m (where m is one of the M chains), the empirical $100(1-\alpha)$ % is initially calculated. The empirical $100(1-\alpha)$ % (e.g., the $(\alpha/2)$ % to $100(1-\alpha)$ % interval) is used in order to calculate a "within-chain" interval estimate for each of the 1 to *M* chains. Using the end points of the within-chain interval for individual chain *m*, the combined observations from all 1 to *M* chains is summed to calculate the percentage of observations from all 1 to *M* chains that lie in the within-chain interval from a single chain *m*. This quantity is denoted here as WCP_m (for within-chain percentage for chain *m*). This is then repeated for all chains, 1 to *M*. The average of the calculated WCP_m for all m (m = 1 to *M*) is then recorded. As *N* (the number of observations in chain *m*) increases, the average $WCP_1 : M$ should approach $100(1-\alpha)$ %, which is the indication of convergence.

In this study, *M* is chosen as 10 chains, while length *N* is varied. The result of this study indicates that it is sufficient to use 1,000 burn-in and 5,000 stored samples to ensure a converged chain. The results of this study are plotted in Figure C.1 for up to 15,000 "stored" samples under five different convergence targets ($\alpha = 10, 5, 2, 1, 0.5$). All chains are visibly approach the expected value by 5,000 stored samples, thus indicating that this combination of burn-in and stored sample numbers is appropriate for this model and application.



Figure C.1: Convergence of MCMC chain of 1,000 burn-in samples and 15,000 stored samples

APPENDIX D CHAPTER 4 OF APPENDIX

The following is a literature review to provide broader context for existing ocean wave observation and monitoring research activity. This appendix chapter is a supporting companion to Chapter 4.

D.1 Background

Characterizing open ocean waters builds upon decades of research, motivated by a wide range of interests, such as weather monitoring [9], rogue wave notification [56], and even evasion of a pirate attack [62]. Various monitoring mechanisms have been employed, from remote sensing instrumentation on satellites to stereophotogrametry, and more recently, computer vision methods integrated in existing data collection methods. This literature review covers the following: (1) use-cases and motivation for monitoring ocean waves, (2) instrumentation used in monitoring waves, and (3) algorithms developed to harness data from instrumentation.

D.2 Use-cases and motivation for monitoring ocean waves

In the current paper, we require a wave monitoring system that could be implemented within the scope of the Arctic ice monitoring framework. Such a system would instantiate a fluid-structures interaction model that could then be used to make predictions on certain mass parameters, and therefore seakeeping abilities, of a ship in operation. While this exact motivation appears to be unique amongst existing literature, there are a number of related research interests that could lend itself to this goal.

Ocean monitoring systems are a way to understand the increased frequency of extreme weather and coastal storms, threats to critical infrastructure, and ecosystems [82]. In essence, monitoring the ocean is a means to monitor both life in the ocean, and how the ocean can impact human lives both in the sea and on land. Specifically, there are a number of applications that have motivated much of the ocean monitoring research, some of which extends as far back as 1939 [64]. Some of the research does not explicitly state an application-based motivation, but implies that the research is generally motivated by marine operations [10, 72].

Most similar to the application that motivates this current work, research from Kos [32] and Gunther [20] specifically cite the desire to integrate wave monitoring systems into navigational and decision support during marine operations. Other relevant researchers have specifically cited a desire to identify single rogue wave events [56] or wave groups [13]. The latter research varies in focus from coastal waves [26], deep water waves [70], shoaling waves [14], and/or short gravity waves [2, 65].

Outside of monitoring waves, much previous research has investigated methods for identifying rigid objects in a maritime scene. This generally involves separating an image into two components: the water scene and the rest of the scene. This can be done by successfully characterizing the water scene through image processing techniques, such that it can be removed. Rigid object detection can be useful as an early warning system for possible collisions [79, 67], and, perhaps more exciting, to evade a possible attack [62].

Another key player that motivates ocean monitoring research is weather forecasting, as the ocean is a valuable indicator of winds and daily weather [9]. This work typically relies on remote sensing from satellites, and is updated on about a daily time scale and may require calibration to compare across differing platforms [85].

Though much of the research focuses on monitoring ocean waves in order to gain knowledge about a scientific or operational concern, some work has been motivated by the desire to generate more realistic waves in a computer graphics setting, for example, in movies, games, or other visual simulation [42, 74, 76].

With the various motivating factors behind studying and monitoring ocean waves, there are numerous dimensions of information to extract about those waves. Much of the research in this field has focused on extracting specifically the directional spectrum of waves [10, 7, 72, 26, 66] or the wavenumber spectra [2]. More recent research has focused on 3-D reconstruction of waves [28, 29], with additional emphasis on extracting significant wave height, wave speed and direction, wind speed and direction, and an understanding of individual waves in time [25, 78, 56, 24]. Other end-forms of data include: wave phase speed [71], slope fields [84], wave climate [85], and wind data [85].

D.3 Ocean wave monitoring instrumentation

A wide range of instrumentation has been used to monitor ocean waves. From 1-D instruments submerged in water (wave buoys) to satellite-borne remote sensing instruments, this is a field of study that has been studied both from direct contact with water, farther away from a satellite in orbit, and nearly any distance in-between.

Wave buoys are standard ground-truth comparisons for all other methods. They are installed throughout the world's oceans, but because they measure one quantity (water surface height) and are installed at specific spatial locations, the data available from this method are inherently limited in number. Thus, there is a tradeoff between spatial resolution and accuracy of information when relying on wave buoy data versus another type of instrumentation. Hence, this is a favored ground-truth method, though not a way to reconstruct the ocean surface with reasonable spatial resolution. Wave buoys have been applied as the ground-truth comparison value for several studies in this overview [85, 36, 43, 47, 29, 76, 12, 24].

There is also ample work on radar systems installed on stable platforms or on aircraft. Researchers have had a growing interest in detecting the spatial and temporal locations of wave groups for safety of offshore platforms, breakwaters, and ships. Through a combination of synthetic aperture radar (SAR) and typical on-board x-band (wavelength ~ 3.2 cm) monitoring, research suggests that it is possible to identify large, extreme wave groups [12]. The use of numerical wave models can be used in conjunction with satellite synthetic aperture radar algorithms and/or radar altimetry to improve the performance of the identification of sea states in extreme situations [36].

Airborne Light Detection and Ranging (LIDAR) is an effective means to obtain an accurate 3-D surface topography of the ocean; however, the cost of flight and data collection does not make it a sustainable or constant means of ocean

145

surface data collection [28, 29]. Satellite-borne remote sensing, such as SAR and radar altimeters, have also been used to infer the sea state [36, 85, 22]. Both significant wave height and wind speed at 10km above the sea surface are standard products that are available for over twenty years of altimeter data. Analysis of this breadth of data suggest that the significant wave height exhibited greater accuracy than the wind speed over the seven missions and time frame analyzed [85]. The repeat cycle for these missions (ERS-1, ERS-2, Envisat, Geosat, GFO, Jason-1, TOPEX) ranged from 3 days to 168 days. Resolution is relatively low, in the kilometers range in deep ocean rather than the meters range. The legacy of altimetry data for significant wave height and concurrent missions makes this technology appealing for wave monitoring applications. With the availability of multiple altimeter platforms (European Remote sensing 1, 2; the envisat, geosat, GFO, Jason-1, and TOPEX), there is a desire to consolidate and compare information across remote sensing datasets of wind speed and wave height. Altimeter data requires offsets to compare across platforms, and such offsets may not be static in time (i.e., the offset for the 1985 to 2008 data may differ from a different set) [85].

In general, very high frequency (VHF) radars are better for wave fields of small heights than high frequency (HF) radars (25-30MHz), and VHF (specifically 47.8 MHz) is recommended for shore, estuarine, and lagoon small scale studies [7]. HF remote sensing of sea state via radar is also common, though these instruments are typically mounted onshore [7]. These systems can reach over ten kilometers offshore and provide high temporal frequency; they are stationary and cannot track waves beyond a specified distance from the shore. This coverage is insufficient for the ships traveling in the Arctic in the proposed application, but its contributions to the general field are worth mentioning here as

part of general background.

Research has also been ongoing in using a marine navigation radar (X-band, around 10 GHz [27]) to measure wave height [83, 78, 47]. Earlier research focused on linear waves with a "relatively simple wave spectra" [83]. Since then, the types of waves and spectra considered have grown in complexity. An available commercial addition, the Wave Monitoring System (WaMoS) connects to existing X-band radar instruments on-board the ship. Resolution for this system has been cited as fine as 5 meters. Though earlier systems were recommended only for lakes or shallow water [78, 47], this system has been been deployed on ships traveling in deeper waters [27], such as the full-scale R/V Melville. Using its on-board X-Band radar, and its IMU, a validation study was conducted in a real sea-state took place near the shore of San Diego, California. This experiment was conducted during a nine-day period in 2013 through an effort by the Office of Naval Research [73]. The comparison of results between the Xband system and airborne LIDAR (considered the ground truth) is reported as having approximately 71% correlation between the LIDAR and X-band wave method in the along- and across- directions of the waves [43]. While this shows promise for future application, alternate options for improved correlation are still needed.

The use of optics in a remote sensing application – specifically, the use of the visible wavelengths – is favorable due to the relatively high spatial resolution, and in a less-remote application, it is also relatively inexpensive (i.e., cameras or video cameras). For example, through the combination of high-quality camera, accurate mapping, fine sampling, long dwell, and large sub-patches, the theoretical dispersion surface could be fit to the observed surface using a simple

algorithm and retrieve the depths when considering only gravity waves [14]. A few variations on the use of optics include polarimetric (filtered optical) techniques, video sequences, and stereophotogrammetry. The holographic method can be used to look at the process of wave generation for specifically deep water, and is relatively easier to process the data than many of its other contemporary methods (stereophotogrammetry, for example) [72]. A study by Zappa in 2008 proves the feasibility of and validates the polarimetric slope sensing technique, such that 2-d slope field of short gravity waves can be accurately recovered without inferring fluid dynamic properties of air or water, and that the recovered water surface still appears realistic [84]. For video application, Stockdon (2000) found large errors associated with wave heights greater than 1m, but introduces the method as an inexpensive, straightforward way to estimate wave phase speed in near shore applications [71]. Another study exampled the feasibility of determining the scale and sea state from a single view, standard video, independent of calibration [69, 70]. This method relies on computer vision techniques, along with prior knowledge of wave dynamics and physics, to infer the wave heights and velocities without any formal calibration stage, making this unique from most other techniques. Resolution is in the correct scale, with the studies ranging from 0.1-0.7 meters per pixel. In another example, as mentioned above, work has been done in analyzing video sequences to separate a marine water scene from rigid objects in the frame, within known size constraints [79, 62, 67]. This is done through pixel-based statistical algorithms. The use of optics is promising though somewhat limiting due to visibility in lowlight (night), or obstructed visibility (such as fog). The logistics of directionality would also have to be considered in order to obtain a panoramic view.

Stereophotogrammetry is the use of a stereo image (two or more images

taken from different positions) in order to estimate the three-dimensional coordinates of points. It requires the identification of common points in each, and most methods rely on a few assumptions about the search path for such common points. The beginnings of stereo pair photographs for analysis of water scenery can be traced as far as 1939 [64], though its more prominent use in determining detailed spectral sea properties begins around 1957 [10]. Counting waves through stereophotogrammetry by hand was commonly used before more major technological advances [26]. Many studies reiterate the labor intensive nature of using stereophotography [66]. In the same study [66], to add to the process, two-dimensions high-pass filtering had to be performed in order to obtain useful information about the short waves; the errors were within 5%. Insensitivity of short wave spectra to wind direction was determined by Banner in 1989 [2]. While one had to be a certified photogrammetrist to work on stereo images (from the Certified Photogrammetrist Program), the start of some manual labor being replaced by digital techniques begins in the early 1990s [65]. By 2004, digital image matching emerged as a way to reduce the intensive labor that is required of manually matching, measuring, and recording stereophotogrammetry [63].

Each sensing instrument type plays to a particular strength; from surveying the literature, it appears that polarimetry is best at capturing short (0.001 - 1m) gravity-capillary and gravity waves, while stereo imaging is favored to capture short- to mid- wavelengths (0.2 - 50m), and X-band radar is most reliable when looking at mid- to long-waves (>10m).

D.4 Ocean wave monitoring algorithms and/or post-processing

Existing data on ocean wind vectors has been used to validate wave height results from video recordings of the ocean [69, 70]. Ocean wind vectors are a parameter closely tied to sea state (a modern version of the Beaufort scale), relating wind speeds to form and height of waves, along with a qualitative description of the sea state. Comparisons to the Beaufort scale have also been made through measurements taken by scatterometers such as RapidScat and its predecessor, QuickSCAT [9].

The use of spectral methods have been proposed by early researchers, but it has not been until relatively recently that the full richness of this method could be exploited due to computational availability [7]. The use of fast fourier tranfer methods by Broche in 1987 [7], for example, was limited to 256 points due to computational limitations. Even as recent as the 2000s, choices for how much data to include in the fast Fourier Transfer (FFT) influence the scope on which this method can be applied. For example, in Spencer [69, 70], a single slice of the video frame is retrieved for analysis due to computational frugality. FFT analysis has also been used to separate water scenery from rigid objects, such as in work interested in identifying approaching hazards to a ship [42, 62].

In more recent years, applications of both image processing, machine learning, and their related sibling computer vision has given rise to a whole new generation of methods that can harness the data from existing instrumentation and collection methods. Of most relevance, the Waves Acquisition Stereo System (WASS) uses a stereo matched image sequence to extrapolate topographic information about a sea surface [3]. This extension of previous work looks into the feasibility and certainty of the WASS system on a moving system, such as onboard a ship [4]. Currently, the stable-camera algorithm is available as an implemented, freely available open-source pipeline [5].

BIBLIOGRAPHY

- [1] T. Acharya and A. K. Ray. Image Transformation. In *Image Processing*, pages 61–78. John Wiley & Sons, Inc., Hoboken, NJ, USA, Sept. 2005.
- [2] M. Banner, I. S. Jones, and J. Trinder. Wavenumber spectra of short gravity waves. *Journal of Fluid Mechanics*, 198:321–344, 1989.
- [3] A. Benetazzo. Measurements of short water waves using stereo matched image sequences. *Coastal Engineering*, 53(12):1013–1032, Dec. 2006.
- [4] A. Benetazzo, F. Barbariol, F. Bergamasco, A. Torsello, S. Carniel, and M. Sclavo. Stereo wave imaging from moving vessels: Practical use and applications. *Coastal Engineering*, 109:114–127, Mar. 2016.
- [5] F. Bergamasco, A. Torsello, M. Sclavo, F. Barbariol, and A. Benetazzo. WASS: An open-source pipeline for 3d stereo reconstruction of ocean waves. *Computers & Geosciences*, 107:28–36, Oct. 2017.
- [6] J. Borge, R. S. Gonzalez, K. Hessner, K. Reichert, and C. G. Soares. Estimation of sea state directional spectra by using marine radar imaging of sea surface. In *Proceedings of the ETCE/OMAE2000 joint conference, New Orleans, ASME*, 2000.
- [7] P. Broche, P. Forget, J. C. De Maistre, J. L. Devenon, and M. Crochet. VHF radar for ocean surface current and sea state remote sensing. *Radio Science*, 22(01):69–75, 1987.
- [8] S. P. Brooks and A. Gelman. General Methods for Monitoring Convergence of Iterative Simulations. *Journal of Computational and Graphical Statistics*, 7(4):434–455, Dec. 1998.
- [9] P. S. Chang, Z. Jeienak, and J. M. Sienkiewicz. Satellite remotely sensed ocean surface vector winds. *Oceanography*, 22(2):194, 2009.
- [10] J. Chase and L. J. Cote. The directional spectrum of a wind generated sea as determined from data obtained by the Stereo Wave Observation Project. Technical report, New York University Bronx School of Engineering and Science, 1957.
- [11] R. Conrad. SMP95: Standard Ship Motion Program User Manual, Dec. 2005.

- [12] H. Dankert. Ocean surface determination from X-band radar-image sequences. *Journal of Geophysical Research*, 109(C4), 2004.
- [13] H. Dankert, J. Horstmann, S. Lehner, and W. Rosenthal. Detection of wave groups in sar images and radar image sequences. *IEEE Transactions on Geoscience and Remote Sensing*, 41(6):1437–1446, June 2003.
- [14] J. Dugan, C. Piotrowski, and J. Williams. Water depth and surface current retrievals from airborne optical measurements of surface gravity wave dispersion. *Journal of Geophysical Research: Oceans*, 106(C8):16903–16915, 2001.
- [15] J. M. Falzarano, S. W. Shaw, and A. W. Troesch. Application of global methods for analyzing dynamical systems to ship rolling motion and capsizing. *International Journal for Bifurcation and Chaos*, 2(1):101–115, 1992.
- [16] D. M. Feit. Forecasting of superstructure icing for alaskan waters. *National Weather Digest*, 12(2):5–10, 1987.
- [17] T. C. Gillmer and B. Johnson. *Introduction to Naval Architecture*. Naval Institute Press, 1982.
- [18] GoPro. Hero 5 Black: User Manual, 2016.
- [19] J. W. Greenert. The United States Navy Arctic Roadmap for 2014 to 2030. Technical report, DTIC Document, 2014.
- [20] H. Gnther, I. Trnkmann, and F. Kluwe. ADOPT DSS-Ocean Environment Modelling for Use in Decision Making support. 2008.
- [21] P. Hansen. *Discrete Inverse Problems: Insight and Algorithms*. Fundamentals of Algorithms. Society for Industrial and Applied Mathematics, 2010.
- [22] K. Hasselmann, R. Raney, W. Plant, W. Alpers, R. Shuchman, D. R. Lyzenga, C. Rufenach, and M. Tucker. Theory of synthetic aperture radar ocean imaging: A MARSEN view. *Journal of Geophysical Research: Oceans*, 90(C3):4659–4686, 1985.
- [23] W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrica*, 57(1):97–109, 1970.
- [24] K. Hathaway and D. Resio. 2-D surface elevation measurements by means of X-band radar. 2006.

- [25] K. Hessner, K. Reichert, and J. Dittmer. Coastal application of a wave monitoring system based on a nautical radar. In *Geoscience and Remote Sensing Symposium*, 1999. IGARSS'99 Proceedings. IEEE 1999 International, volume 1, pages 500–502. IEEE, 1999.
- [26] L. Holthuijsen. Observations of the directional distribution of oceanwave energy in fetch-limited conditions. *Journal of Physical Oceanography*, 13(2):191–207, 1983.
- [27] W. Huang and E. W. Gill. Ocean Remote Sensing Using X-Band Shipborne Nautical RadarApplications in Eastern Canada. In *Coastal Ocean Observing Systems*, pages 248–264. Elsevier, 2015.
- [28] P. A. Hwang, D. W. Wang, E. J. Walsh, W. B. Krabill, and R. N. Swift. Airborne measurements of the wavenumber spectra of ocean surface waves. Part I: Spectral slope and dimensionless spectral coefficient. *Journal of Physical Oceanography*, 30(11):2753–2767, 2000.
- [29] P. A. Hwang, D. W. Wang, E. J. Walsh, W. B. Krabill, and R. N. Swift. Airborne measurements of the wavenumber spectra of ocean surface waves. Part II: Directional distribution. *Journal of Physical Oceanography*, 30(11):2768–2787, 2000.
- [30] JCGM. Evaluation of measurement data guide to the expression of uncertainty in measurement. Technical Report JCGM 100:2008 GUM 1995 with minor corrections, Joint Committee for Guides in Metrology, Bureau International des Poids Mesures (BIPM), Svres, France, 2008.
- [31] B. Kinsman. *Wind waves: their generation and propagation on the ocean surface.* Courier Corporation, 1965.
- [32] S. Kos, R. Ive, and D. Bri. VTS improvements introduction of wave monitoring system as a navigational sensor. In *International Symposium on Electronics in Transport*, 2012.
- [33] A. Kulyakhtin, O. Shipilova, B. Libby, and S. Løset. Full-scale 3d cfd simulation of spray impingement on a vessel produced by ship-wave interaction. In 21st IAHR International Symposium on Ice, Dalian, China, 2012.
- [34] D. Lee, S. Y. Hong, and G. J. Lee. Theoretical and experimental study on dynamic behavior of a damaged ship in waves. *Ocean Engineering*, 34:21– 311, 2007.

- [35] E. Lee, C. R. Weil, and A. Fullerton. Experimental Results for the Calm Water Resistance of the Generic Prismatic Planing Hull (GPPH). Technical Report NSWCCD-80-TR-2017/015, Naval Surface Warfare Center Carderock Division, May 2017.
- [36] X. M. Li, S. Lehner, and M. X. He. Ocean wave measurements based on satellite synthetic aperture radar (SAR) and numerical wave model (WAM) data extreme sea state and cross sea analysis. *International Journal of Remote Sensing*, 29(21):6403–6416, Nov. 2008.
- [37] Y. C. Lin and C. J. Earls. Multi-parameter stochastic inversion for first and second moment mass properties of model-scale ship with topside ice accumulation. *Submitted to Applied Ocean Research (April 29, 2018)*.
- [38] Y. C. Lin and C. J. Earls. Validation experiment of a single-view image sequence algorithm to identify scale and sea state characteristics. *Submitted to Applied Ocean Research (May 13, 2018)*.
- [39] Y. C. Lin and C. J. Earls. Stochastic inversion framework to monitor evolving mass properties of a ship at sea during Arctic opeations. In *The 30th American Towing Tank Conference Proceedings*, West Bethesda, Maryland, Oct. 2017.
- [40] Y. C. Lin, C. J. Earls, J. T. Park, and T. C. Smith. Stochastic inversion for the roll gyradius second moment mass property in ships at full-scale and model-scale. *Applied Ocean Research*, 63:24–35, Feb. 2017.
- [41] W. A. Link and R. J. Barker. *Bayesian Inference with ecological examples*. Academic Press, 2010.
- [42] G. A. Mastin, P. A. Watterberg, and J. F. Mareda. Fourier synthesis of ocean scenes. *IEEE Computer graphics and Applications*, 7(3):16–23, 1987.
- [43] C. Merrill, J. Geiser, W. Pfitsch, T. Fu, E. Terrill, T. dePaolo, T. Cook, L. Lenain, S. Richer, and A. Fullerton. Simultaneous Ship Motions and Ocean Waves Measured in the Time Domain. Nov. 2014.
- [44] L. M. Minnick, L. W. Hanyok, H. A. Tomaszek, M. P. Melendez, C. R. Turner, and J. T. V. B. C. C. B. Park. Model-scale experiment of the seakeeping performance for r/v melville, model 5720. Technical Report NSWCCD-50-TR-2012/035, Naval Surface Warfare Center Carderock Division, 2012.

- [45] U. D. Nielsen, I. M. V. Andersen, and J. Koning. Comparisons of means for estimating sea states from an advancing large container ship. In *The* 12th International Symposium on Practical Design of Ships and Other Floating Structures, Changwon City, Korea, October 2013.
- [46] U. D. Nielsen and D. Stredulinsky. Sea state estimation from an advancing ship - a comparative study using trial data. *Applied Ocean Research*, 34:33– 44, 2012.
- [47] J. Nieto Borge, G. R. Rodrguez, K. Hessner, and P. I. Gonzlez. Inversion of marine radar images for surface wave analysis. *Journal of Atmospheric and Oceanic Technology*, 21(8):1291–1300, 2004.
- [48] W. L. Oberkampf and M. F. Barone. Measures of agreement between computation and experiment: Validation metrics. *Journal of Computational Physics*, 217(1):5–36, Sept. 2006.
- [49] W. L. Oberkampf and C. J. Roy. *Verification and validation in scientific computing*. Cambridge University Press, 2010.
- [50] W. L. Oberkampf and T. G. Trucano. Verification and validation in computational fluid dynamics. *Progress in Aerospace Sciences*, page 64, 2002.
- [51] C. Octopus. /f180 series data sheet, 2012.
- [52] J. E. Overland. Prediction of vessel icing for near-freezing sea temperatures. *Weather and forecasting*, 5(1):62–77, 1990.
- [53] J. Park, C. R. Turner, and M. P. Melendez. Physical Properties and Roll Decay with Uncertainty Estimates for DTMB Model 5720, 23rd Scale R/V Melville. Technical Report NSWCCD-80-TR-2015/000, Naval Surface Warfare Center Carderock Division, 2016.
- [54] S. Punzi. Roll-Decay Testing of Model 5748 Data Description. Technical report, NSWCCD.
- [55] Python Software Foundation. Python language reference, version 2.7, 2001.
- [56] K. Reichert, K. Hessner, J. Dannenberg, and I. Traenkmann. X-band radar as a tool to determine spectral and single wave properties. In 25th International Conference on Offshore Mechanics and Arctic Engineering, pages 683– 688. American Society of Mechanical Engineers, 2006.

- [57] C. J. Roy and W. L. Oberkampf. A comprehensive framework for verification, validation, and uncertainty quantification in scientific computing. *Computer Methods in Applied Mechanics and Engineering*, 200(25-28):2131– 2144, June 2011.
- [58] C. C. Ryerson. Superstructure spray and ice accretion on a large US Coast Guard cutter. *Atmospheric research*, 36(3):321–337, 1995.
- [59] C. C. Ryerson. Icing management for coast guard assets. Technical report, Cold Regions Research and Engineering Laboratory, Hanover, NH, 2013.
- [60] C. C. Ryerson and A. J. Gow. Ship Superstructure Icing: Crystalline and Physical Properties. Technical report, DTIC Document, 2000.
- [61] C. C. Ryerson and P. D. Longo. Ship Superstructure Icing: Data Collection and Instrument Performance on USCGC MIDGETT Research Cruise. Technical report, DTIC Document, 1992.
- [62] J. G. Sanderson, M. K. Teal, and T. J. Ellis. Characterisation of a complex maritime scene using Fourier space analysis to identify small craft. 1999.
- [63] F. Santel, W. Linder, and C. Heipke. Stereoscopic 3d-image sequence analysis of sea surfaces. In *Proceedings of the ISPRS Commission V Symposium*, volume 35, pages 708–712, 2004.
- [64] A. Schumacher. *Stereophotogrammetrische Wellenaufnahmen*, volume 7. W. de Gruyter, 1939.
- [65] O. Shemdin and H. Minh Tran. Measuring short surface waves with stereophotography. *Photogrammetric engineering and remote sensing*, 58(3):311–316, 1992.
- [66] O. H. Shemdin, H. M. Tran, and S. Wu. Directional measurement of short ocean waves with stereophotography. *Journal of Geophysical Research: Oceans*, 93(C11):13891–13901, 1988.
- [67] A. A. W. Smith, M. K. Teal, and P. Voles. The statistical characterization of the sea for the segmentation of maritime images. In *Video/Image Processing and Multimedia Communications*, 2003. 4th EURASIP Conference focused on, volume 2, pages 489–494. IEEE, 2003.
- [68] F. G. W. Smith. *The seas in motion*. Thomas Y. Crowell, Company, 1973.

- [69] L. Spencer and M. Shah. Water Video Analysis. 2004.
- [70] L. Spencer, M. Shah, and R. Guha. Determining scale and sea state from water video. *IEEE Transactions on Image Processing*, 15(6):1525–1535, June 2006.
- [71] H. F. Stockdon and R. A. Holman. Estimation of wave phase speed and nearshore bathymetry from video imagery. *Journal of Geophysical Research: Oceans*, 105(C9):22015–22033, 2000.
- [72] Y. Sugimori. A study of the application of the holographic method to the determination of the directional spectrum of ocean waves. In *Deep Sea Research and Oceanographic Abstracts*, volume 22, pages 339–350. Elsevier, 1975.
- [73] E. J. Terrill. X-band Observations of Waves, Algorithm Development, and Validation High Resolution Wave-Air-Sea Interaction DRI. Technical report, Scripps Institute of Oceanography, 2012.
- [74] J. Tessendorf. Simulating ocean water. *Simulating nature: realistic and interactive techniques. SIGGRAPH*, 1(2):5, 2001.
- [75] J. M. T. Thompson, R. C. T. Rainey, and M. S. Soiman. Ship stability criteria based on chaotic transients from incursive fractals. *Philosophical Transactions of the Royal Society*, A, 332:149–167, 1990.
- [76] S. Thon and D. Ghazanfarpour. Ocean waves synthesis and animation using real world information. *Computers & Graphics*, 26(1):99–108, 2002.
- [77] E. C. Tupper. Introduction to Naval Architecture, fourth edition. Elsevier, 2004.
- [78] J. Vogelzang, K. Boogaard, K. Reichert, and K. Hessner. Wave height measurements with navigation radar. *International Archives of Photogrammetry and Remote Sensing*, 33(B7/4; PART 7):1652–1659, 2000.
- [79] P. Voles, M. Teal, and J. Sanderson. Target identification in a complex maritime scene. 1999.
- [80] WAFO-group. WAFO A Matlab Toolbox for Analysis of Random Waves and Loads - A Tutorial. Math. Stat., Center for Math. Sci., Lund Univ., Lund, Sweden, 2017.

- [81] G. B. Whitham. Non-linear dispersion of water waves. *Journal of Fluid Mechanics*, 27(02):399, Feb. 1967.
- [82] Z. Willis. National Ocean Observing Systems in a Global Context. In *Coastal Ocean Observing Systems*, pages 11–25. Elsevier, 2015.
- [83] I. R. Young, W. Rosenthal, and F. Ziemer. A three-dimensional analysis of marine radar images for the determination of ocean wave directionality and surface currents. *Journal of Geophysical Research: Oceans*, 90(C1):1049– 1059, 1985.
- [84] C. J. Zappa, M. L. Banner, H. Schultz, A. Corrada-Emmanuel, L. B. Wolff, and J. Yalcin. Retrieval of short ocean wave slope using polarimetric imaging. *Measurement Science and Technology*, 19(5):055503, May 2008.
- [85] S. Zieger, J. Vinoth, and I. R. Young. Joint calibration of multiplatform altimeter measurements of wind speed and wave height over the past 20 years. *Journal of Atmospheric and Oceanic Technology*, 26(12):2549–2564, 2009.