

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Freshwater Processes in the Upper Ocean

Permalink

<https://escholarship.org/uc/item/4gw8g7n6>

Author

Hoffman, Lauren

Publication Date

2023

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Freshwater Processes in the Upper Ocean

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Chemical Engineering

by

Lauren A. Hoffman

Committee in charge:

Matthew Mazloff, Chair
Jinhye Bae, Co-Chair
Donata Giglio
Sarah Gille
Shyue Ping Ong
Justin Opatkiewicz

2023

Copyright

Lauren A. Hoffman, 2023

All rights reserved.

The Dissertation of Lauren A. Hoffman is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2023

EPIGRAPH

It isn't knowledge that's making trouble,
but the uses it's put to.

Kurt Vonnegut

TABLE OF CONTENTS

Dissertation Approval Page	iii
Epigraph	iv
Table of Contents	v
List of Figures	viii
List of Tables	xiii
Acknowledgements	xiv
Vita	xvi
Abstract of the Dissertation	xvii
Introduction	1
Chapter 1 Ocean Surface Salinity Response to Atmospheric River Precipitation in the California Current System	6
1.1 Summary	6
1.2 Introduction	7
1.3 Background	8
1.3.1 Salinity variability in the California Current System	8
1.3.2 Salinity response to precipitation	9
1.3.3 Atmospheric rivers in the California Current System	10
1.3.4 Impacts of salinity on global moisture distribution	11
1.4 Observational Data and Model	12
1.4.1 Instrument Accuracy	13
1.4.2 ERA5	13
1.4.3 SIO-R1 AR Catalog	14
1.4.4 CUGN Spray Line 66.7	15
1.4.5 MBARI M1 Mooring	15
1.4.6 MITgcm 1D Model	16
1.5 Methods	16
1.5.1 Seasonal Time Scale	16
1.5.2 Event Studies	20
1.6 Results	24
1.6.1 Seasonal Response	24
1.6.2 Event-Based Response	28
1.7 Discussion	36
1.7.1 Seasonal Response	37
1.7.2 Event-Based Response	38

1.8	Conclusion	40
1.9	Acknowledgments	41
1.10	Data Statement	42
Chapter 2	Machine learning for daily forecasts of Arctic sea-ice motion: an attribution assessment of model predictive skill.	44
2.1	Summary	44
2.2	Introduction	45
2.3	Background	47
2.4	Data	50
2.4.1	Sea-Ice Velocity: Polar Pathfinder Version 4 Daily Sea Ice Motion vectors (PP)	51
2.4.2	Sea-Ice Concentration: Nimbus-7 SMMR and DMSP SSM/I-SSMIS Passive Microwave Data	51
2.4.3	Wind Velocity: Japanese 55-year Reanalysis derived for ocean-ice models (JRA55-do)	52
2.4.4	Bathymetric Depth: International Bathymetric Chart of the Arctic Ocean (IBCAO)	53
2.5	Methods	53
2.5.1	Model Inputs	53
2.5.2	Model Setup	54
2.5.3	Model Evaluation	57
2.5.4	Model Comparison	59
2.5.5	Analysis of Inputs	59
2.6	Results	60
2.6.1	Model Performance	60
2.6.2	Linear Regression Parameters: Relationship Between Sea-Ice Motion and Input Parameters	68
2.6.3	Attribution assessment of model predictive skill	72
2.7	Conclusions	82
2.7.1	A CNN can make skillful predictions of sea-ice motion on one-day time scales.	82
2.7.2	Model predictive skill and discrepancies between model performances are linked to various properties related to sea-ice motion.	83
2.7.3	Wind velocity plays the largest role in predicting ice velocity.	84
2.8	Acknowledgements	85
2.9	Data Statement	86
Chapter 3	Evaluating the robustness of explainable machine learning (XML) methods for application in regression predictions of Arctic sea-ice motion.	87
3.1	Summary	87
3.2	Introduction	88
3.3	Data	90
3.4	Methods	91

3.4.1	Model Setup	91
3.4.2	Sensitivity Analyses	92
3.5	Results	97
3.5.1	Localized Sensitivity Studies	97
3.5.2	Domain-Integrated Sensitivity Studies	101
3.6	Discussion and Conclusions	105
Chapter 4	Explainable machine learning (XML) for evaluating trends in the drivers of variability in Arctic sea-ice dynamics.	110
4.1	Summary	110
4.2	Introduction	111
4.3	Data	113
4.4	Methods	113
4.4.1	The Models: Linear Regression (LR) and Convolutional Neural Network (CNN)	113
4.4.2	Trends in the ‘Relevance’ of the Inputs	114
4.5	Results	115
4.5.1	Overall LR Parameters and Trends	115
4.5.2	Spatial Variability of LR Parameters Compared to Inputs	116
4.5.3	Spatial Trends in LR Parameters and Inputs	117
4.6	Discussion	119
4.6.1	How is the wind factor related to ice concentration?	119
4.6.2	Future Work: LRP with CNN	124
4.7	Conclusions	126
Chapter 5	Conclusions and Future Work	128
5.1	Surface salinity response to precipitation from atmospheric rivers in the California Current System.	128
5.2	Assessing the viability of machine learning (ML) methods for predicting and understanding Arctic sea-ice dynamics.	130
5.3	Characterizing changes in the relationship between ice drift and wind.	133
5.4	Interdisciplinarity	136
Appendix A	Supplemental Material for “Ocean Surface Salinity Response to Atmospheric River Precipitation in the California Current System”	137
Appendix B	Supplemental Material for “Machine learning for daily forecasts of Arctic sea-ice motion: an attribution assessment of model predictive skill.”	143
Appendix C	Supplementary Material for “Evaluating the robustness of explainable machine learning (XML) methods for application in regression predictions of Arctic sea-ice motion.”	147
Bibliography	165

LIST OF FIGURES

Figure 1.1.	(a) Fraction of rain events with precipitation greater than 5 mm day^{-1} that are also ARs; and (b) fraction of total precipitation that comes from ARs, within the region of the CCS. Events included occur between September and March for the years 2007-2019.	11
Figure 1.2.	Observed vs. modeled March-minus-September salinity differences (psu) color coded by cumulative rainfall (cm) for the years 2008–2019. The solid black line represents the linear regression of observed to modeled salinity data for all rain rates, plotted with 99% confidence	19
Figure 1.3.	Time series showing salinity (psu) for MITgcm one-dimensional model runs (red, solid) and MBARI M1 Mooring (red, dashed) at 1 m depth, compared to ERA5 rain rate (mm day^{-1}) (blue) from September through March in 2015–2018 (a–d). The black dashed line represents the initial . . .	24
Figure 1.4.	(a) Climatological annual cycle and (b) multi-year time series of salinity anomaly as a function of offshore distance at 10 m depth as measured by the CUGN Spray underwater glider on line 66.7. (c,d) Salinity anomaly averaged over offshore distances from 0–50 km (red) and daily	25
Figure 1.5.	(a-c) Cumulative rainfall (cm) and (d-f) number of AR events as a function of salinity (psu) difference between March and September for the years 2008–2019 at (a,d) offshore, (b,e) onshore, and (c,f) coastal locations. Panels a–c include CUGN Spray line 66.7 observations (blue)	27
Figure 1.6.	Composite time series of (a) six-hourly rain (mm, red) and daily cumulative rain from day -3 to day n (mm, blue), (b) wind speed (m s^{-1}) with a six-hour moving mean, and (c) salinity difference (psu) between relative day n and relative day 0 for 85 rain events occurring at the MBARI M1	29
Figure 1.7.	Normalized salinity anomaly in the upper 55 m of the ocean for the four-day one-dimensional MITgcm runs and for wind speeds from 2 to 16 m s^{-1} and maximum rain rates from 2 to 8 mm h^{-1} . Each contour plot is divided by the absolute value of the maximum salinity anomaly for the	31
Figure 1.8.	Results from the MITgcm experiments using idealized environmental forcing in which the peak rain rate and the wind speed are varied. (a) Peak magnitude of ΔS , ΔS_{max} , as a function of rain rate for five different wind speeds; (b) ΔS_{max} as a function of wind speed for different rain rates	32

Figure 1.9.	Same as Fig. 1.8 a & b, zoomed in to enhance view of results from event case studies (colored circles). The colored circles show model output from event case studies, with the colors representing wind speed and rain rate, respectively. The black dotted line represents the salinity difference	33
Figure 1.10.	Results from case studies for three AR events in the CCS. (a–c) Time series of rain rate (mm h^{-1} , blue) and wind speed (m s^{-1} , red) that was used as model forcing from ERA5 at the coastal location; (d–f) time series showing the salinity difference (ΔS , psu) from the first time step at	35
Figure 2.1.	Maps showing (a) spatial divisions (Greenland Sea, Eastern Arctic, Central Arctic, and Baffin Bay), (b) bathymetric depth (m) (note logarithmic scaling), and (c) the distance from coast (km). Spatial divisions are based on overall performance of the CNN model and the difference	48
Figure 2.2.	Schematic of the convolutional neural network (CNN) used in this study for predicting present-day sea-ice velocity components, $u_{i,t}$ & $v_{i,t}$ (outputs), from present-day wind velocity, $u_{a,t}$ & $v_{a,t}$, previous-day sea-ice velocity, $u_{i,t-1}$ & $v_{i,t-1}$, and previous-day sea-ice concentration, $c_{i,t-1}$	56
Figure 2.3.	(a–c) Mapped correlation for predictions of sea-ice velocity made by the (a) PS, (b) LR, and (c) CNN models. (d–f) The difference in correlation between models: (d) $\text{corr}_{\text{CNN}} - \text{corr}_{\text{PS}}$, (e) $\text{corr}_{\text{LR}} - \text{corr}_{\text{PS}}$, and (f) $\text{corr}_{\text{CNN}} - \text{corr}_{\text{LR}}$. The gray regions in d–f represent locations where	63
Figure 2.4.	PDFs for model performances compared to their distance from the coast (a–c) and bathymetric depth (d–f), with (a & d) for the CNN and (b & e) for the LR. (c & f) The difference ($\text{corr}_{\text{CNN}} - \text{corr}_{\text{LR}}$) between the correlation of the two models. Gray shading (c & f) represents correlation	64
Figure 2.5.	(a–d) Ensemble mean monthly correlation for the prediction of sea-ice velocity by three different models: (a) all models, (b) persistence, (c) linear regression, and (d) CNN. (e–g) The difference between the correlation of the (e) CNN and PS, (f) LR and PS, and (g) CNN and LR models.	67
Figure 2.6.	Correlation of the CNN (a–c) and the difference between CNN and LR correlation (d–f) as a function of various properties related to sea-ice motion (wind speed, u_a (a & d); sea-ice speed, u_i (b & e); and sea-ice concentration, c_i (c & f). The correlation is calculated with	69
Figure 2.7.	(a–c) Magnitude of the normalized linear regression coefficient for the relationship between sea-ice velocity components and input parameters (a, wind speed, A ; b, sea-ice speed, B ; c, sea-ice concentration, C) normalized to the maximum of a–c. (d–e) Mean angle of (d) wind speed	71

Figure 2.8.	PDFs for LR correlation (a–c) and the difference between the correlation of the CNN & LR correlation (d–f) compared to the LR coefficient magnitudes for (a & d) wind speed, LRu_a ; (b & e) sea-ice speed, LRu_i ; and (c & f) sea-ice concentration, LRc_i	72
Figure 2.9.	Overall mean of the performance metrics, ((a) $corr_{CNN}$, and (b) $corr_{CNN} - corr_{LR}$), and properties related to ice motion ((c) wind velocity, u_a ($m s^{-1}$); (d) ice velocity, u_i ($m s^{-1}$); (e) ice concentration, c_i ; (f) bathymetric depth, d (m); (j) distance from coast, d_c (km); and the LR coefficients	74
Figure 2.10.	Ensemble mean of the correlation between the model performance metrics (circles for $corr_{CNN}$ in a & c; triangles for $corr_{CNN} - corr_{LR}$ in b & d) and the various properties related to ice motion within each of the spatial divisions	75
Figure 3.1.	Schematic of the convolutional neural network (CNN) used in this study for predicting present-day sea-ice velocity components (outputs)	93
Figure 3.2.	Locations used in integrated analysis for global LRP implementation (blue dots). Regions of the Arctic are labeled for reference during discussion. . .	98
Figure 3.3.	Spatial mean and standard deviation of the “relevance” scores for the covariance, perturbation, and LRP sensitivity methods	102
Figure 3.4.	Results from localized sensitivity studies for each of the inputs at location 9, indicated by the red dot. The columns represent each of the different inputs: (a), (d), and (g), wind speed, u_a ; (b), (e), and (h), ice speed, u_i ; and (c), (f), and (i), ice concentration, c_i).	103
Figure 3.5.	(a–i) Probability density of the relevance of each localized sensitivity study as a function of the distance from the point of analysis for each of the inputs and for all locations. The columns represent each of the different inputs: (a), (d), and (g), wind speed, u_a ; (b), (e), and (h), ice speed	104
Figure 3.6.	Results from global sensitivity studies for each of the inputs. The columns represent each of the different inputs: (a), (e), and (i), wind speed, u_a ; (b), (f), and (j), ice speed, u_i ; (c), (g), and (k), ice concentration, c_i ; and (d), (h), and (l), the maximum parameter at each location.	106
Figure 4.1.	Monthly LR parameters and model inputs from 1989–2021. LR parameters for (a) wind speed, LR, u_a , (b) ice speed, LR, u_i , and (c) ice concentration, LR, c_i . Monthly (d) wind speed, u_a , ice speed, u_i , and (f) ice concentration. In (a)–(f), black lines and error bars represent	119

Figure 4.2.	Map of seasonal mean for (a–d) the LR parameters for wind speed, LR, u_a , (e–h) the ice speed, u_i , and (i–l) the ice concentration, c_i from 1989–2021. Each column represents a different season (i.e. January–March, April–June, July–September, October–December).	120
Figure 4.3.	Map of seasonal trend in (a–d) the LR parameters for wind speed, LR, u_a , (e–h) the ice speed, u_i , and (i–l) the ice concentration, c_i from 1989–2021. Each column represents a different season (i.e. January–March, April–June, July–September, October–December).	121
Figure A.1.	Example MITgcm simulation used for four-day sensitivity tests with idealized forcing from day -2 to day 2. (a) Time series of wind (m s^{-1} , red) and rain (mm hr^{-1} , blue) forcing. (b) Salinity (psu) response in the upper 20 m. The thin black line indicates the depth of the fresh lens, D_L	138
Figure A.2.	Example of external forcing from ERA5 and initial conditions from Spray for the case study with an AR event starting on October 17, 2016. (a) rain rate (mm hr^{-1}) and wind speed (m s^{-1}), (b) short and longwave radiation (W m^{-2}), (c) specific humidity (kg kg^{-1}), and atmospheric	139
Figure A.3.	Comparison of MITgcm vs. GOTM model output of maximum salinity difference, ΔS_{max} , for event-based studies using conditions in the tropics and wind speeds of 4–10 m s^{-1} . The black line represents a linear regression between the two sets of model results, with slope and r^2	140
Figure A.4.	Collection of 85 rain events used in the composite analysis at the MBARI M1 Mooring location from day -3 to day 7. (a) Daily cumulative rain (mm) from ERA5, (b) hourly wind speed (m s^{-1}) with a six hour moving mean from ERA5 and (c) surface salinity anomaly (psu) from	141
Figure A.5.	Monte Carlo simulation to find uncertainty on the parameters A and b in the linear regression equation relating rain rate and precipitation to change in salinity, $\Delta S_{max} = AR_{max}U^b$. In (a) data from the one-dimensional model output (black) for various wind and rain conditions and the	142
Figure B.1.	Uncertainty metrics for the Polar Pathfinder Sea Ice Motion Vector, Version 4 data set. Monthly (blue) and annual (red) mean (a) error of sea-ice velocity with error bars (b) error of sea-ice velocity without error bars, and (c) standard deviation of the error in sea-ice velocity	144
Figure B.2.	Ensemble mean (a) wind factor and (b) turning angle for the relationship between wind velocity and sea-ice velocity calculated from the linear regression equation: $u_{i,t}^* = Au_{a,t}^* + B$. Data was not standardized in this case because here we represent the wind factor as the ratio	146

Figure C.1.	Results from localized sensitivity studies for each of the inputs at location 1, indicated by the red dot. The columns represent each of the different inputs: (a), (d), and (g), wind speed, u_a ; (b), (e), and (h), ice speed, u_i ; and (c), (f), and (i), ice concentration, c_i	148
Figure C.2.	Same as Fig. C.1, but for location 2.	149
Figure C.3.	Same as Fig. C.1, but for location 3.	150
Figure C.4.	Same as Fig. C.1, but for location 4.	151
Figure C.5.	Same as Fig. C.1, but for location 5.	152
Figure C.6.	Same as Fig. C.1, but for location 6.	153
Figure C.7.	Same as Fig. C.1, but for location 7.	154
Figure C.8.	Same as Fig. C.1, but for location 8.	155
Figure C.9.	Same as Fig. C.1, but for location 9.	156
Figure C.10.	Same as Fig. C.1, but for location 10.	157
Figure C.11.	Same as Fig. C.1, but for location 11.	158
Figure C.12.	Same as Fig. C.1, but for location 12.	159
Figure C.13.	Same as Fig. C.1, but for location 13.	160
Figure C.14.	Same as Fig. C.1, but for location 14.	161
Figure C.15.	Same as Fig. C.1, but for location 15.	162
Figure C.16.	Same as Fig. C.1, but for location 16.	163
Figure C.17.	Same as Fig. C.1, but for location 17.	164

LIST OF TABLES

Table 1.1.	Model parameters for (a) seasonal (b) event sensitivity and (c) event case studies.	43
Table 2.1.	Overall correlation and skill between observations and predictions of sea-ice velocity for four different models.	60
Table 3.1.	Details about training data for the CNN and LR models, and model architecture (i.e. parameters and hyperparameters) for the CNN.	109
Table B.1.	Comparison of model data and architecture to Zhai and Bitz (2021).	145
Table B.2.	Comparison of model performance to Zhai and Bitz (2021); calculated from the temporal evaluations (Fig. 5a) for consistency.	146

ACKNOWLEDGEMENTS

Thanks to the waves and the rocks, fluids and solids, for keeping me both flexible and grounded. I'd like to acknowledge a number of people to whom I am sincerely grateful. Some of you helped me keep my solid footing and gave me the encouragement I needed to keep going. Others pushed the limits of my worldview, showing me that I could bend without breaking and grow stronger from it. All of you helped me see the resilience of the human spirit and what can be achieved with drive, dedication, and curiosity.

The culmination of this work would not have been possible without Matt Mazloff, Sarah Gille, and Donata Giglio. Your brilliant insights and supportive guidance were invaluable to this work and to my growth as a scientist. Thank you to my thesis committee, Jinhye Bae, Shyue Ping Ong, and Justin Opatkiewicz for investing time and energy on a topic slightly outside of the realm of a typical chemical engineering curriculum. Special thanks to Kayli Matsuyoshi, Aniruddh Varadarajan, Cecilia Bitz, and Patrick Heimbach for your wonderful contributions to this work. I would also be remiss in not mentioning all of you who took the time to meet with me and provide helpful perspectives on machine learning and sea ice: thank you Libby Barnes, Kirsten Mayer, Will Chapman, and Lorenzo Zampieri. Additionally, I am thankful for everyone who has sat through one of my presentations and given me an active audience to share my science with.

This dissertation has benefited immensely from oceanographic and atmospheric observational and modeling communities, including the California Underwater Glider Network, the MBARI moorings, the Center for Western Weather and Water Extremes (CW3E) Atmospheric River Catalog, the MITgcm, the ERA5 and JRA55 reanalysis products, and the National Snow and Ice Data Center satellite products.

The work in this dissertation was funded by an Office of Naval Research Grant N00014-20-1-2772. Additional support was received from the Center for Western Weather and Water Extremes (CW3E).

I also would like to acknowledge my appreciation for David Miller who hired me as a TA

for fluid dynamics seven different times, and who was there to support me during my PhD before I started my research journey. Additionally, to the students I taught, thank you for helping me forever overcome my fear of public speaking - now it's one of my favorite things about what I do!

I could not have taken this journey without the support of my family and friends. Thank you Lily, thank you Alice, and thank you Truc for being there for me and giving me space to share myself. Last, but not least, I'd like to thank my cats Fourier and Laplace for the comfort and entertainment. Thank you for sitting on my lap every day, and for limiting yourselves to minor distractions... for the most part.

Chapter 1, in full, is a reprint of the material as it appears in *Journal of Physical Oceanography* 2023. This chapter is published as: Hoffman, L., Mazloff, M. R., Gille, S.T., Giglio, D., and Varadarajan, A. (2022). Ocean Surface Salinity Response to Atmospheric River Precipitation in the California Current System. *Journal of Physical Oceanography*, 52(8): 1867–1885. ©American Meteorological Society. Used with permission. The dissertation author was the primary investigator and author of this paper.

Chapter 2 is a reprint of the material as it may appear in *Artificial Intelligence for Earth Systems*, 2023. We cite this work as: Hoffman, L., Mazloff, M. R., Gille, S.T., Giglio, D., Bitz, C. M., Heimbach, P., and Matsuyoshi, K. (2022). Machine learning for daily forecasts of arctic sea-ice motion: an attribution assessment of model predictive skill. *Artificial Intelligence for Earth Systems*. ©American Meteorological Society. Used with permission. This preliminary version has been accepted for publication in *Artificial Intelligence for the Earth Systems* and may be fully cited. The final typeset copyedited article will replace the EOR when it is published. The dissertation author was the primary investigator and author of this paper.

VITA

- 2016 B. S. in Mechanical Engineering *magna cum laude*, University of San Diego, San Diego
- 2016-2018 Lead Engineer, Primo Wind, Inc., San Diego
- 2019 M. S. in Chemical Engineering, University of California San Diego
- 2018-2020 Graduate Teaching Assistant, University of California San Diego
- 2020-2021 Instructor, University of California San Diego, Academic Connections
- 2023 Ph. D. in Chemical Engineering, University of California San Diego

PUBLICATIONS

Hoffman, L., M. Mazloff, S.T. Gille, D. Giglio, P. Heimbach, C. Bitz, K. Matsuyoshi, “Machine learning for daily forecasts of Arctic sea-ice motion: an attribution assessment of model predictive skill”, , 2023.

Hoffman, L., M. Mazloff, S.T. Gille, D. Giglio, A. Varadarajan., “Ocean Salinity Response to Atmospheric River Precipitation Events in the California Current System”, *Journal of Physical Oceanography*, 52, 1867-1885, 2023.

FIELDS OF STUDY

Major Field: Chemical Engineering

Studies in Physical Oceanography

Studies in Machine Learning

ABSTRACT OF THE DISSERTATION

Freshwater Processes in the Upper Ocean

by

Lauren A. Hoffman

Doctor of Philosophy in Chemical Engineering

University of California San Diego, 2023

Matthew Mazloff, Chair

Jinhye Bae, Co-Chair

Freshwater exchanges between the ocean–ice–atmosphere system play a crucial role in the global climate system. This study provides an analysis of the local impact of freshwater fluxes both off the coast of California and in the Arctic. Studies are carried out using observations and numerical and statistical models. We show that freshwater exchanges between the ocean and atmosphere in the form of precipitation from atmospheric rivers (ARs) over the ocean in the California Current System (CCS) have impacts on the surface ocean salinity on event and seasonal timescales. In the upper ocean, precipitation from ARs can produce long-lasting layers of freshwater, the extent of which are dependent on atmospheric forcing from precipitation

and wind. We conclude that upper ocean salinity changes due to ARs are within the limits of detectability of ocean instruments.

We also examine the extent to which wind acts as a driving force for ice motion in the Arctic. To accomplish this, we build a sequence of machine learning (ML) models that make one-day predictions of present-day zonal and meridional sea-ice velocity components from inputs of present-day wind velocity, previous-day sea-ice velocity, and previous-day sea-ice concentration. We analyze the performance of these models, and implement explainable machine learning (XML) methods to understand how they are making their predictions. One of these methods, layerwise relevance propagation (LRP), was developed for ML models that make classification rather than regression predictions. This study is the first known implementation of a global LRP for a regression problem in geosciences. We therefore provide a comparative study of several different XML methods to bolster trustworthiness in the use of LRP for this particular application. A convolutional neural network (CNN) has improved performance compared conventional persistence (PS) and linear regression (LR) models. Outputs from local LRP studies are shown to be consistent with other XML methods. However global implementations of LRP are highly sensitive to choices made during processing. We analyze the coefficients of the LR model to understand the relationship between ice motion and wind speed. We show that the ice is becoming more responsive to wind forcing, and link this to decreasing ice concentration.

Introduction

Distribution of freshwater in the upper ocean plays an important role in the global climate system. The exchange of freshwater between the ocean and atmosphere is a key component of the global freshwater cycle (Solomon et al., 2021). Freshwater fluctuations in the upper ocean are a balance between sources (precipitation, river runoff, ice melt), sinks (evaporation, ice growth), and redistribution (horizontal advection, vertical mixing). These processes are closely linked to atmospheric variability, which impact the distribution of freshwater through transport of mass, momentum, and heat. In addition to direct exchanges of freshwater between the ocean and atmosphere, wind also plays a large role in upper ocean freshwater budgets. Processes linked to freshwater fluxes have both local and remote impacts (Wijesekera and Boyd, 2001).

Locally, freshwater exchanges in the ocean influence upper-ocean salinity and density gradients (i.e. stratification). Surface stratification controls vertical mixing and impacts the development of the surface mixed layer, which can cause mixed-layer entrainment that further modulates the oceanic response to atmospheric forcing (Yamaguchi and Suga, 2019). For example, the formation of long-lasting freshwater layers at the surface can inhibit turbulent vertical mixing and decrease exchanges between the mixed layer and the thermocline (Schmitt, 2008). This can lead to the formation of diurnal warm layers (Webster et al., 1996), enhanced surface currents (Wijesekera et al., 1999), the suppression of near-surface turbulent dissipation below lenses (Smyth et al., 1997), and may provide unexpected regional variation of internal wave energy propagation, dissipation, and mixing in the thermocline (Schmitt, 2008).

On a global scale, freshwater fluxes in the upper ocean play a role in large-scale oceanic and atmospheric circulation. For example, freshwater modulation due to the melting of ice in the

Antarctic is related to deep water formation, which impacts large-scale overturning circulations (Goosse and Fichefet, 1999). This thermohaline circulation (i.e. circulation caused by density gradients related to variations in temperature and salinity) is linked to variations in global climate on decadal to millennial timescales (Schmitt, 1995). Additionally, Sévellec et al. (2017) used an optimal flux perturbation framework to show that on multi-decadal timescales anomalous fluxes of freshwater in the Arctic that result from a declining sea ice extent are the dominant mechanism in driving weakening of the Atlantic Meridional Overturning Circulation (AMOC).

Freshwater in the Arctic is also modulated by sea ice advection around the Arctic and export through the Fram Strait. Export of sea ice through the Fram Strait is the largest dynamic sink of freshwater in the Arctic, and has been increasing in response to changes in the climate state of the Arctic (Solomon et al., 2021). Rapid changes in the Arctic climate system have led to a regime shift to what has been dubbed the “New Arctic” (Jeffries et al., 2013). These changes include a reduction in the sea ice extent (Stroeve and Notz, 2018), thinning of the ice cover (Kwok, 2018), warming and freshening of the Arctic Ocean (Timmermans and Marshall, 2020), regionally enhanced oceanic and atmospheric mixing, and enhanced ocean–ice–atmosphere coupling (Polyakov et al., 2020). These changes have a large impact on the variability of freshwater in the Arctic and around the globe.

Understanding the impacts and redistribution of freshwater is important because the global freshwater cycle is predicted to amplify as the vapor-carrying capacity of the atmosphere increases from expected changes in climate (Yu et al., 2020). Changes in the global water cycle are cited as some of the most societally relevant aspects of climate change (Lagerloef et al., 2010; SPURS-2 Planning Group, 2015; Yu et al., 2020). This work explores the implications of atmospheric river (AR) precipitation and Arctic sea-ice dynamics on freshwater budgets and air-sea interactions in the upper ocean. Analyses are carried out using both observations and modeling. However, the bulk of the analyses are approached from a modeling perspective, with implementations of both dynamical and statistical models.

The first part of this work analyzes the impact of AR precipitation on upper ocean

salinity in the California Current System (CCS). Salinity variability in the CCS has historically been attributed to intrinsic ocean dynamics, including alongshore advection from the California Current (Lynn and Simpson, 1987; Schneider et al., 2005), upwelling (Aquad et al., 2011), and riverine runoff (Kudela and Chavez, 2004; Johnson et al., 1999). The salinity response to precipitation in the CCS has received little attention to date, particularly in contrast to studies in the tropics that show a significant impact of upper ocean salinity on precipitation (Boutin et al., 2014; Clayson et al., 2019; Drucker and Riser, 2014; Drushka et al., 2016, 2019). This work investigates the role of local atmospheric forcing in the form of precipitation from ARs in influencing upper ocean salinity. Atmospheric rivers are examples of extreme weather events that impact the distribution of heat and moisture between the ocean and atmosphere, both locally and globally. ARs are associated with extreme precipitation on land and over the ocean, and account for up to 82% of total rainfall in the CCS. Climate projections indicate that the moisture content and frequency of extreme AR events and seasons are expected to increase in response to a warming climate (Dettinger, 2011; Payne et al., 2020; Shields and Kiehl, 2016). Understanding how ARs impact surface ocean salinity, mixing, and stratification is important for understanding their impacts on air-sea exchanges at different timescales and at different locations. This work uses a combination of observations and a one-dimensional configuration of the MITgcm (Massachusetts Institution of Technology General Circulation Model; a dynamical model) to analyze the surface salinity response to ARs on seasonal and event timescales. The response of the ocean to freshwater input is a function of precipitation, wind, background stratification, heat flux, and vertical mixing (Drushka et al., 2016).

The remainder of this work focuses on using ML to predict and understand sea-ice dynamics in the Arctic. The analysis begins by assessing the viability of using a neural network as a surrogate for the dynamical component of ice in a numerical model that provides nowcasting of the state of Arctic sea ice. ML models are promising for this application because they do not need to resolve the complex physics inherent to numerical models of sea-ice dynamics, and therefore have improved computational efficiency and cost (Hunke et al., 2020). Statistical

models of ice dynamics have historically been based on a linear relationship between wind speed and ice speed, where ice motion is related to wind through a speed reduction factor (the wind factor) and a turning angle (Thorndike and Colony, 1982). This relationship describes up to 70% of the variability in ice motion in the central Arctic for short time scales, but has decreased success in coastal regions (Thorndike and Colony, 1982; Kimura and Wakatsuchi, 2000; Hibler, 1979; Maeda et al., 2020; Kwok et al., 2013). In comparison to linear regression (LR), ML models in the form of neural networks incorporate non-linear relationships between the inputs into their predictions, and are able to capture important spatial information. In this study we develop a convolutional neural network (CNN) that can make skillful one-day predictions of ice motion. ML models are trained to predict present-day zonal and meridional ice velocity components from inputs of present-day wind velocity, previous-day ice velocity, and previous-day ice concentration. These predictors are chosen based on the momentum equation for sea-ice, where scale analyses show that the dominant drivers of ice motion on short time scales are wind and ocean currents, and internal ice stresses (Hibler, 1979; Lepparanta, 2011; Olason and Notz, 2014; Feltham, 2008). We analyze the overall performance of the CNN in comparison to classical statistical models. We also analyze the spatial and temporal distributions in performance and how they are linked to the variability of various properties related to ice motion.

In addition to assessing the predictive skill of ML models, this study also investigates the “black box” nature of ML models and analyzes what information the models are using to make predictions. This methodology, known as explainable ML (XML), provides insight into how the model output (i.e. sea-ice velocity) is linked to the various inputs and how those relationships are changing in time. We begin with a methods-based approach, where we investigate several different XML techniques. We perform a comparative analysis of outputs from standard methods of interpretability with newer methods that are able to provide more information. In particular, we aim to understand the layerwise relevance propagation (LRP) explainability method because of the sheer quantity of information that it provides. This study is the first known application of a domain-integrated implementation of LRP to a regression problem in the geosciences. Therefore

we compare it to other XML methods to provide context and build trust in the method.

Ultimately, we aim to use output from XML to understand the relationship between ice speed and wind speed and how it is changing in response to the changing ice state. Increases in ice drift speeds have been attributed to ice thinning and reduction of multiyear sea ice coverage (Rampal et al., 2009; Spreen et al., 2011; Zhang et al., 2012; Kwok et al., 2013; Carmack et al., 2015; Tandon et al., 2017; Docquier et al., 2017). Reductions in sea-ice extent lead to a regime where more of the ice is in free drift, and reductions in ice thickness lead to changes in ice mechanics, both of which increase the responsiveness of ice motion to winds. Increases in the wind factor between ice and wind speed have been cited for many regions throughout the Arctic for all seasons (Spreen et al., 2011; Maeda et al., 2020). In this study we use ML in the form of a linear regression (LR) model to confirm the increase in the relationship between ice and wind speed, and to understand the mechanisms driving it.

Chapter 1

Ocean Surface Salinity Response to Atmospheric River Precipitation in the California Current System

1.1 Summary

Atmospheric rivers (ARs) result in precipitation over land and ocean. Rainfall on the ocean can generate a buoyant layer of fresh water that impacts exchanges between the surface and the mixed layer. These “fresh lenses” are important for weather and climate because they may impact the ocean stratification at all timescales. Here we use in situ ocean data, co-located with AR events, and a one-dimensional configuration of a general circulation model, to investigate the impact of AR precipitation on surface ocean salinity in the California Current System (CCS) on seasonal and event-based time scales. We find that at coastal and onshore locations the CCS freshens through the rainy season due to AR events, and years with higher AR activity are associated with a stronger freshening signal. On shorter time scales, model simulations suggest that events characteristic of CCS ARs can produce salinity changes that are detectable by ocean instruments (≥ 0.01 psu). Here, the surface salinity change depends linearly on rain rate and inversely on wind speed. Higher wind speeds ($U > 8 \text{ m s}^{-1}$) induce mixing, distributing freshwater inputs to depths greater than 20 m. Lower wind speeds ($U \leq 8 \text{ m s}^{-1}$) allow freshwater lenses to remain at the surface. Results suggest that local precipitation is important in setting

the freshwater seasonal cycle of the CCS and that the formation of freshwater lenses should be considered for identifying impacts of atmospheric variability on the upper ocean in the CCS on weather event time scales.¹

1.2 Introduction

Freshwater inputs from rainfall can have variable impacts on surface ocean salinity. Of particular significance is the impact on upper-ocean stratification, which has been shown to limit the penetration depth of wind mixing and thus the vertical distribution of atmospheric fluxes (Schmitt, 2008; Chaudhuri et al., 2021; Thompson et al., 2019). This has larger implications for intensification of the global water cycle (SPURS-2 Planning Group, 2015; Yu et al., 2020). The relative importance of factors that are known to impact the ocean's response to freshwater inputs is not well characterized, especially in the subtropics where studies are limited. Atmospheric Rivers (ARs) are narrow, elongated plumes of strong poleward water vapor transport known to produce large amounts of precipitation over the ocean and land in the California Current System (CCS) (Ralph and Dettinger, 2012; Ralph et al., 2013). The impact of ARs on surface ocean salinity has received minimal attention to date. Previously, global seasonal salinity variations in the upper ocean have been attributed to runoff (in coastal regions), advection in the ocean, as well as evaporation and precipitation (Yu, 2011). Ren and Riser (2009) found that among these, in the subarctic regions of the Northeast Pacific (45°N - 50°N), precipitation was the largest contributor. However, they did not address the California Current System, where variations in salinity have been linked to variations in anomalous advection along the trajectories of the California Current, the Inshore Current, and the California Undercurrent on seasonal (Lynn and Simpson, 1987), interannual, and decadal (Schneider et al., 2005) timescales. Therefore to date, seasonal variations of salinity within the CCS have mainly been attributed to advection (Lynn and Simpson, 1987; Schneider et al., 2005). Here we hypothesize that local precipitation in the

¹This chapter is published as Hoffman, L., Mazloff, M. R., Gille, S.T., Giglio, D., and Varadarajan, A. (2022). Ocean Surface Salinity Response to Atmospheric River Precipitation in the California Current System. *Journal of Physical Oceanography*, 52(8): 1867–1885. ©American Meteorological Society. Used with permission.

CCS (including ARs) provides a significant contribution to seasonal freshening. Additionally, we hypothesize that precipitation from ARs impacts the surface ocean on shorter time scales, and may be detectable by oceanographic salinity sensors in some conditions.

This study uses a combination of observations and modeling with the aim of understanding the surface salinity response to ARs in the California Current System by characterizing (i) the ocean salinity response to precipitation over the duration of the wet season; and (ii) the role of rain rate and wind speed in driving changes in upper-ocean salinity and stratification for characteristic AR events on event time scales.

1.3 Background

1.3.1 Salinity variability in the California Current System

Surface salinity variability in the CCS is typically attributed to alongshore advection from the California Current (Lynn and Simpson, 1987; Schneider et al., 2005). Situated 150–1300 km offshore, the California Current transports cool, fresh, nutrient-rich water southward. Within the coastal zone (0–150 km) there is a poleward flow of warm, saline, low-oxygen subtropical waters from the California Inshore Countercurrent (IC) (Bograd et al., 2001; Lynn and Simpson, 1987). At the surface (upper 50 m), the IC has seasonality, with a poleward flow occurring in the winter and fall, and an equatorward flow in the spring and summer (Lynn and Simpson, 1987; Rudnick et al., 2017b). Salinity increases toward the coast, implying that an increase in offshore flow would result in an increase in salinity offshore (Rudnick et al., 2017b). Additionally, in a study of the temperature and salinity extremes found in the CCS beginning in 2017, Ren and Rudnick (2021) concluded that the positive salinity anomaly was a result of advection and that different source waters were found in the California Current from 2017-2019. During the summer, the increased salinity at the coast is enhanced due to coastal upwelling of cold, saline waters from depth (Aquad et al., 2011). Riverine runoff has been linked to salinity decreases off the coast of central California (Kudela and Chavez, 2004; Johnson et al., 1999). While, as noted

in the introduction, salinity variability in the CCS has previously been attributed to intrinsic ocean dynamics (Lynn and Simpson, 1987; Schneider et al., 2005; Auad et al., 2011; Kudela and Chavez, 2004; Johnson et al., 1999), atmospheric forcing such as local surface freshwater flux may also influence surface salinity and is investigated here.

1.3.2 Salinity response to precipitation

The response of the ocean to freshwater input is a function of rainfall, wind, background stratification, heat flux, and vertical velocity in the upper ocean (Drushka et al., 2016). Rainfall forms stably stratified upper-ocean layers, with lenses of fresher water of $O(1 \text{ m to } 10 \text{ m})$ thick. Changes in these freshwater lenses are driven by the interaction between buoyancy and shear forces; they can persist from minutes to hours depending on factors such as wind-driven surface mixing, lateral advection, convective overturning during nighttime cooling, and internal and surface waves (Brainerd and Gregg, 1997; Drushka et al., 2019; Price, 1979; Tomczak, 1995; Wijesekera et al., 1999). While most fresh layers disperse within a few hours, in some cases fresh layers have been shown to persist for tens of hours (Walesby et al., 2015). Long-lasting freshwater layers can inhibit turbulent vertical mixing and decrease exchanges between the mixed layer and the thermocline (Schmitt, 2008). This can lead to the formation of diurnal warm layers (Webster et al., 1996), enhanced surface currents (Wijesekera et al., 1999), and the suppression of near-surface turbulent dissipation below lenses (Smyth et al., 1997). In addition, fresh lenses may provide unexpected regional variation of internal wave energy propagation, dissipation, and mixing in the thermocline (Schmitt, 2008). While this work pertains to freshwater lenses rather than barrier layers (Soloviev et al., 2015), it is interesting to note that de Boyer Montégut et al. (2007) identified the presence of unexplained barrier layers off the California coast at $25\text{--}45^\circ$ latitude. This study may explain the mechanisms behind this previously unexplained phenomenon.

While the ocean salinity response to precipitation in the CCS has received little attention to date, there is a growing pool of research on the ocean's response to freshwater input in

the tropics, as experiments involving Surface Salinity Profilers (SSP) provide high-resolution measurements near the surface. Results from a SSP deployed in the western tropical Pacific in December 2011 indicate that the vertical salinity difference between 0.26 m and 0.11 m depth has a cubic dependence on rain rate, and is inversely proportional to wind speed (Asher et al., 2014). Other studies have shown a linear relationship between the vertical salinity gradient and maximum rain rate (Boutin et al., 2014; Clayson et al., 2019; Drucker and Riser, 2014; Drushka et al., 2016, 2019). However, wind speed was not factored into all of these studies. In the cases where wind was taken into account, results from a one-dimensional general ocean turbulence model (GOTM) and measurements made in the Intertropical Convergence Zone (ITCZ) in the eastern tropical Pacific during the second Salinity Processes in the Upper-ocean Regional Study (SPURS-2) showed the maximum difference in salinity between 1-5 m depth and the surface to be inversely proportional to wind speed (Drushka et al., 2016, 2019). In this study, we focus on the subtropics, where studies to date have been limited.

1.3.3 Atmospheric rivers in the California Current System

ARs account for a substantial amount of the global water transport, especially at mid-latitudes where they can supply more than 90% of meridional transport of atmospheric water vapor (Ralph and Dettinger, 2012; Zhou and Newell, 1998). ARs are characterized by high atmospheric water vapor content and heavy winds. Because they are associated with extreme precipitation on land and over the ocean, especially in coastal regions (Ralph and Dettinger, 2012; Ralph et al., 2013), ARs often cause devastating flooding and play a large role in the global distribution of moisture and drought (Ralph and Dettinger, 2011). ARs can occur in families consisting of several (typically 2–6) consecutive ARs (Fish et al., 2019), contributing to the accumulation of precipitation in the upper ocean and on land. The AR that extends from Hawaii to the US West Coast carries moisture across the eastern Pacific to the coast of California. Off the coast of Monterey Bay in the CCS, 30-48% of precipitation events greater than 5 mm day^{-1} occur during ARs, which are responsible for up to 82% of total rainfall in the CCS, as seen along

California Cooperative Oceanic Fisheries Investigations (CalCOFI) line 66.7 in Fig. 1.1, and as indicated by Guan and Waliser (2015). Argo profiles indicate large-scale upper ocean freshening on average from December to February in areas of the Pacific that receive frequent AR-associated rainfall (Giglio et al., 2020). Implications of AR events for upper-ocean stratification and salinity are important, especially as climate projections indicate that the moisture content of ARs and the frequency of extreme AR events and storm seasons are expected to increase as a result of a warming climate (Dettinger, 2011; Payne et al., 2020; Shields and Kiehl, 2016).

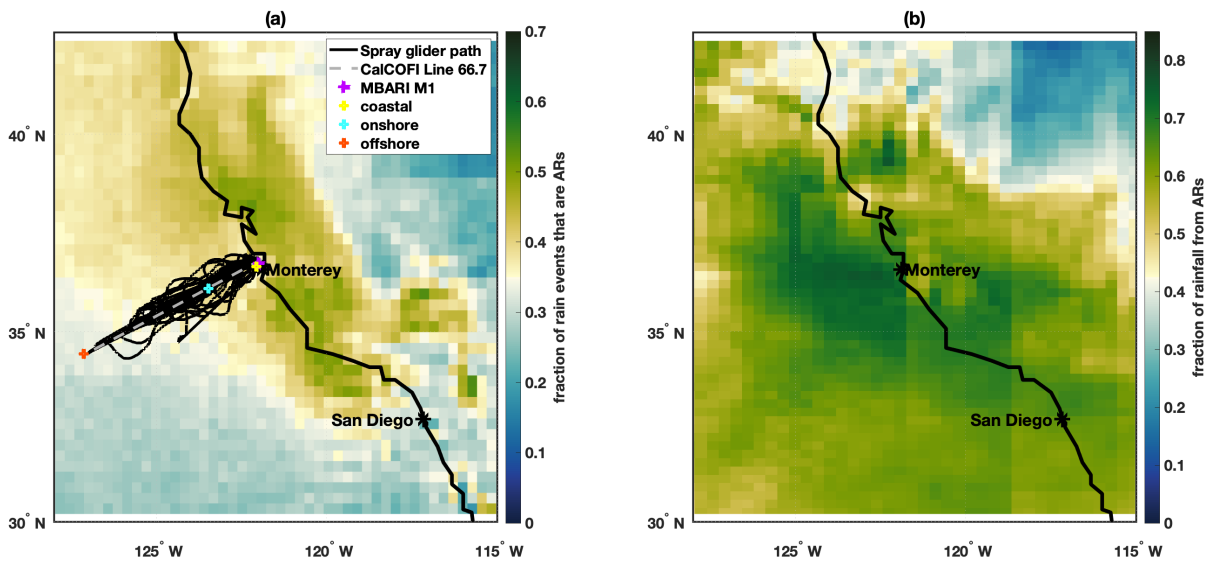


Figure 1.1. (a) Fraction of rain events with precipitation greater than 5 mm day⁻¹ that are also ARs; and (b) fraction of total precipitation that comes from ARs, within the region of the CCS. Events included occur between September and March for the years 2007-2019. Also depicted is the trajectory traveled by CUGN Spray glider along CalCOFI line 66.7, the location of the MBARI M1 mooring (purple) and the coastal (yellow), onshore (cyan), and offshore (red) locations that were used during model analyses. The gray dashed line represents CalCOFI line 66.7 off the coast of Monterey, CA.

1.3.4 Impacts of salinity on global moisture distribution

Changes in surface salinity have broad implications for the distribution of moisture and the Earth’s water cycle. For example, a reduction in sea surface salinity due to precipitation is hypothesized to lead to a positive feedback in which the formation of buoyant freshwater

layers reduces vertical mixing in the upper ocean, which then contributes to increased SST, and in turn leads to a further increase in atmospheric convection and precipitation (SPURS-2 Planning Group, 2015). In contrast, Williams et al. (2006) used climate modeling to show that freshwater lenses formed from an intensified hydrological cycle could produce a basin-scale negative sea surface temperature feedback to anthropogenic human climate change. These nuances make understanding the vertical upper-ocean salinity gradient important for improving air-sea coupling in models (McCulloch et al., 2012) and understanding the role of upper ocean stratification in a changing climate. Boutin et al. (2013) also suggested that the impact of precipitation on salinity stratification should be taken into account when assimilating satellite data under rainy conditions. Furthermore, the Clausius-Clapeyron relationship shows a strong, non-linear dependence of water vapor pressure on temperature. With this relation, a rise in temperature of about 1°C leads to a 7% increase in vapor pressure, which causes changes in the water cycle as the vapor-carrying capacity of the atmosphere increases (Schmitt, 2008). These changes will impact the global distribution of rainfall and drought, which is one of the most societally relevant aspects of climate change (SPURS-2 Planning Group, 2015; Yu et al., 2020).

1.4 Observational Data and Model

A combination of observations and modeling are used to determine the seasonal and event-based response of ocean salinity to rain events within the CCS (30°N - 42.5°N , 128°W - 115°W). Here the region is divided into three subdomains based on the distance from shore: coastal (0-50 km), onshore (50-150 km) and offshore (150-550 km). The distance ranges are chosen based on the location of California Undercurrent (strongest around 70 km offshore), the California Inshore Countercurrent (strongest around 150 km offshore), and the California Current (strongest at 200–300 km offshore) as they fall along CalCOFi line 66.7 (Rudnick et al., 2017b). The subdomains include data collected along the Spray glider line, and their bounds, perpendicular to the coast, are indicated by three colored markers in Fig. 1.1. Model initialization

and forcing data are taken from observations and reanalysis fields at three coordinate locations (36.67°N, 122.06°W; 36.11°N, 123.47°W; and 34.43°N, 127.13°W, which are 30 km, 150 km, and 550 km offshore from Monterey Bay, respectively) within the three subdomains (coastal, onshore, and offshore). Figure 1.1 shows these locations and indicates the location of the Spray glider path along CalCOFi line 66.7 and the Monterey Bay Aquarium Research Institute (MBARI) M1 mooring location.

1.4.1 Instrument Accuracy

The accuracy specification for conductivity, temperature, depth (CTD) instruments in measuring salinity is equivalent to 0.003 psu. However, this value is defined in a clean, well-mixed calibration bath and does not take into account effects of in situ ocean measurements. For example, the dynamic effects of moving instruments are known to increase errors in CTD measurements to 0.02-2.0 psu (Seabird Scientific, 2016). This is consistent with observation errors for in situ salinity data that are found to be typically on the order of ± 0.01 psu after post-processing for quality control (Vinogradova et al., 2019; Delcroix et al., 2005). These values are similar to the 0.01 psu accuracy reported in Argo salinity measurements after delayed-mode adjustments (Wong et al., 2020). Here, we use 0.01 psu as the threshold for a detectable salinity change.

1.4.2 ERA5

The ERA5 dataset is produced using a 4D-Var data assimilation of the European Centre for Medium-Range Weather Forecasts (ECMWF) Integrated Forecast System (IFS) by combining a vast number of historical observations into global estimates. Covering the Earth on a 31 km (0.28128°) grid and resolving the atmosphere using 137 levels from the surface to 80 km height, the ERA5 dataset provides hourly estimates of a number of surface ocean and atmospheric variables from 1979 to present (Hersbach et al., 2020). In an analysis of the performance of five state-of-the-art global reanalyses in comparison to in situ data, ERA5 surface winds were found

to have the best agreement with observed variability on daily and interannual time scales (Ramon et al., 2019). The ERA5 dataset showed significant improvements in precipitation estimates compared to ERA-Interim, with the caveat that biases still remained in the southeastern United States and on the North American western coast (Tarek et al., 2020). Additionally, reanalysis products (including the ERA5) showed the best agreement with precipitation measurements made by local ground stations in a comparison of a collection of satellite, reanalysis, and gauge measurements from the Frequent Rainfall Observations on GridS (FROGS) dataset for two case studies (California and Portugal) of extreme AR events (Ramos et al., 2021). However, the ERA5 often underestimated heavy precipitation compared to gauge measurements, with a mean absolute percent error of 68% (Ramos et al., 2021).

In this study, the ERA5 reanalysis dataset (Muñoz Sabater, 2019) is used to characterize atmospheric conditions, i.e. atmospheric temperature, T_a (K); zonal and meridional wind speed, U_Z and U_M (m s^{-1}); downwelling longwave radiation and shortwave radiation, I_L and I_S (W m^{-2}); specific humidity, SpH (kg kg^{-1}); evaporation minus precipitation, EmP (m s^{-1}); and rain rate, R (m s^{-1}). This study uses hourly data at the surface within the CCS from 2007-2019 to match the date range of the dataset for the Spray glider along line 66.7.

1.4.3 SIO-R1 AR Catalog

The Scripps Institution of Oceanography (SIO)-generated AR catalog, the SIO-R1 AR catalog (Gershunov, 2017), provides a record of AR activity on the North American West Coast (20.0° - 60.0° N, 160° - 100° W). The dataset indicates whether or not an AR was detected (0 or 1) for each 6-hourly time step on a 2.5° resolution spatial grid (Gershunov et al., 2017). Here, this catalog is used to investigate the fraction of events with rainfall exceeding 5 mm day^{-1} that are associated with ARs (Fig. 1.1), as well as the total number of AR events during the rainy season each year. Here we define the AR as ‘detected’ if there is an AR in the grid cell or neighboring grid cell. To quantify rain events, we use ERA5 precipitation estimates at the AR locations.

1.4.4 CUGN Spray Line 66.7

The California Underwater Glider Network (CUGN) provides continuous sampling along CalCOFI line 66.7 by one Spray glider at a time (Rudnick, 2016). The glider travels from Monterey Bay to a distance about 500 km offshore, vertically profiling in a sawtooth pattern. Each cycle to 500 m depth and back to the surface covers 3 km of horizontal distance and takes roughly 2.75 h. The quality controlled Spray glider dataset provides temperature and salinity observations from the glider ascent phase at discrete 10 m vertical levels, with the shallowest measurements available at 10 m depth (Davis et al., 2008). Finer resolution (raw) data are available, but performing quality control at depths shallower than 10 m is beyond the scope of this study. Salinity collected by the Spray glider is reported in practical salinity units (psu). Data are available from April 2007 through present (Rudnick et al., 2017b). Here glider data are used to characterize the ocean's salinity response to atmospheric precipitation on seasonal time scales and to initialize model runs (as described in sections 1.5.1 and 1.5.2). Spray glider data allow us to investigate precipitation impacts on salinity at larger spatial scales over the CCS. One limitation of the Spray dataset for this study is that the temporal response of the upper-ocean salinity to precipitation is not fully captured at a particular location due to the fact that the glider is neither a Lagrangian nor an Eulerian platform and is travelling cross-shore.

1.4.5 MBARI M1 Mooring

The MBARI M1 mooring (Chavez, 2015) measures continuously at one location. Therefore in comparison to Spray it has the disadvantage of conveying no spatial information, but the advantage of not aliasing spatial variability into temporal fluctuations. Here we use surface measurements (nominal depth of 1 m) of ocean salinity at a location 20 km offshore of Monterey Bay (36.75° N, -122.0° W; purple marker in Figure 1.1) from 2007 - 2019. This dataset is used to investigate the seasonal response of salinity to precipitation, to compare to model output, and to make event composites.

1.4.6 MITgcm 1D Model

In this study, a one-dimensional configuration of the MITgcm (Adcroft et al., 2018), with vertical transport equations for momentum and heat, is used to run both seasonal (September - March) and event-based simulations (four-day sensitivity studies and nine-day case studies) aimed at characterizing the ocean’s response to precipitation from ARs on different time scales. The MITgcm uses the non-local K Profile Parameterization (KPP) vertical mixing scheme of Large et al. (1994) with a standard configuration as listed in Adcroft et al. (2018). Turbulent heat fluxes are computed in the model using methods from Large and Pond (1982). Details of model setup for each experimental run (seasonal, event sensitivity, and event case studies) are provided in Table 1.1 and in the sections that follow.

1.5 Methods

1.5.1 Seasonal Time Scale

Observational Methods

The seasonal response of ocean salinity is first investigated by looking at the MBARI M1 mooring surface (1 m) salinity measurements from 2015-2018, which are compared with model output from simulations run at the mooring location. Model forcing and initialization are discussed further in section 1.5.1. This is followed by analysis of the annual and interannual (2008 through 2019) salinity anomaly from the Spray glider along line 66.7 in the CCS. As part of this analysis we assess a one-dimensional salinity budget at a location 15 km offshore along the glider path using the hypothesis that changes in salinity within the water column will be fully explained by $E - P$ in the form of an equation,

$$\frac{d}{dt} \left(\frac{\int_0^Z S dz}{Z} \right) = \frac{(E - P)S_{\text{ref}}}{Z} \quad (1.1)$$

Here we ignore advection and diffusion and calculate the amount of precipitation required to produce the rain-year salinity anomaly over a depth, Z , in the limiting case where evaporation, E (from ERA5), and rain, P , are the only contributing factors.

Additionally, over the rain-year from September through March, cumulative precipitation is calculated from ERA5 and compared with change in salinity at 10 m depth from the Spray glider along line 66.7 in coastal, onshore, and offshore regions. Glider offshore distance is calculated by comparing Spray glider data for latitude and longitude at given time steps with the initial coordinate location 5 km offshore. Salinity data are binned monthly and into coastal, onshore and offshore subdomains for each year, and averaged over each bin. Changes in salinity from September (start of the rain-year) to March are calculated for each year from the averages of the binned values. Along the line 66.7 glider path, ERA5 precipitation data are extracted at the fixed locations used to represent the coastal, onshore, and offshore regions, respectively (Fig. 1.1). ERA5 data from each location are binned by month to calculate cumulative monthly precipitation, from which cumulative precipitation is calculated from September through March, to be compared with change in salinity. Uncertainties for salinity and rainfall between September and March are computed by calculating the standard error of the mean in each bin and then propagating errors through the calculations to produce cumulative rainfall or salinity differences.

Model Setup

The seasonal, one-dimensional MITgcm model is run over a period of 213 days (September 1–April 1) with a 0.5 h time step. Atmospheric forcing is applied daily and taken from ERA5 daily mean (longwave and shortwave radiation, zonal and meridional winds, atmospheric temperature, and specific humidity) and daily cumulative (precipitation) values. Forcing is applied for three different locations representing the coastal, onshore, and offshore subdomains. Initial conditions are taken to be temperature and salinity depth profiles, interpolated to 0.5 m intervals, from the Spray glider dataset along line 66.7, which provides measurements at 10 m intervals. The shallowest Spray measurements are at 10 m, so T and S between 0 and 10 m

are set to the 10-m values, under the assumption of a well-mixed surface layer with constant T and S in the upper 10 m. Profiles of T and S are binned by month and by offshore distance for each year. Initial profiles are set as the calculated average profiles in September for each year (2008-2019) and offshore distance regime. When no data are available for September in a given year/distance bin, the T and S profiles from October are used as initial conditions. This is the case for 2008 (coastal bin), 2012 (coastal and onshore bins), and 2017 (coastal bin). The model is run for the upper 140 m of the water column, using 280 vertical levels with 0.5 m spacing. The depth of 140 m was chosen to allow ample room for the downward propagation of the salinity response, as even for cases of high wind speeds, the salinity response to freshwater input was not found to propagate below 120 m depth. These model parameters are also listed in Table 1.1.

Model Validation

The use of a one-dimensional model will allow for analysis without the influence of ocean processes such as horizontal advection, upwelling, and runoff, thus isolating the impact of rainfall and wind speed on upper-ocean salinity changes. We validate the model for long-term studies by comparing the observed and modeled March-minus-September salinity differences for all rain rates over the years 2008–2019 (Fig. 1.2). To do this, the methods discussed in section 1.5.1 for Spray glider data are applied to model output. A linear regression of observed to modeled salinity difference finds a slope of 1.25 with an r^2 value of 0.52, which is statistically significant at the 99% level. Figure 1.2 also shows that a 1:1 ratio between observed and modeled data falls within the 99% prediction interval (green shading) and is close to the upper bound of the 99% confidence interval (blue shading) for the linear fit. Here the prediction interval represents the estimated range of a future observation, while the confidence interval represents the range of values for the linear regression slope and indicates how well this slope has been determined. Higher cumulative rainfall in Fig. 1.2 typically corresponds to a larger rainy-season decrease in salinity, as seen in the gradient of the color-coded data points, where large negative salinity differences (salinity decrease) are dark blue (high cumulative rain), and large positive

salinity differences (salinity increase) are tan (low cumulative rain). Spray salinity differences tend to be larger than modeled differences, indicated by the slope being slightly large than one (i.e. for every 1 unit change in modeled salinity difference, Spray measures a change of 1.25 units). This difference in slope could be indicative of the model not including horizontal advection, upwelling, or runoff.




Figure 1.2. Observed vs. modeled March-minus-September salinity differences (psu) color coded by cumulative rainfall (cm) for the years 2008–2019. The solid black line represents the linear regression of observed to modeled salinity data for all rain rates, plotted with 99% confidence (blue shading) and prediction (green shading) intervals. The slope and r^2 value for the fit are indicated in the legend. The black dotted line indicates the 1:1 relationship. Data are included from coastal, onshore, and offshore locations. With 27 data points, linear regression coefficients are statistically different from zero at the 99% confidence level if $r^2 > 0.24$; our results exceed this threshold.

1.5.2 Event Studies

Observational Methods

To assess the salinity response to precipitation on an event basis, we analyze ERA5 precipitation at the location of the MBARI M1 Mooring surface salinity measurements. Event composites are created by averaging rainfall, wind speed and salinity from 85 heavy rain events as a function of time relative to the start date, described below. Events are included if daily cumulative precipitation is greater than a threshold of 5 mm and there has not been another rain event of this size within 10 days prior to the event start date. Events are defined to start (day 0) on the first date with rainfall exceeding the threshold. For the MBARI M1 mooring, events are chosen within a date range from January 2007 through March 2019. Composite analysis is not carried out using data from the Spray glider. While the decrease in salinity in response to precipitation is visible for a few glider events (not shown), the motion of the Spray glider makes composites too difficult to compute in a consistent way.

Model Setup, Sensitivity Studies

Event-based sensitivity studies are run in the one-dimensional configuration of the MITgcm for four-day periods to study the impact of AR events on the formation of freshwater lenses. Atmospheric forcing is applied every minute with the 60-s time steps linearly interpolated from hourly ERA5 fields. In order to isolate the impact of wind speed on surface mixing, values for radiation (I_L and I_S), specific humidity (SpH), and air temperature T_a are kept constant and set as the calculated average value of the ERA5 dataset over five coastal AR events from October 2016 – February 2017. Characteristic precipitation, wind speed, and event duration are defined based on commonly occurring conditions for AR events, as noted in the statistical distribution of different conditions for composited AR events from Table 2 in Ralph et al. (2013). Precipitation is applied as a 12-hour long Gaussian pulse (defined by the full width of the Gaussian at one tenth of the peak) with maximum rain rate ($R = 0, 2, 3, 4, 5, \text{ and } 8 \text{ mm h}^{-1}$) occurring during the 48th hour, preceded and followed by a period of zero rainfall. The Gaussian pulse was chosen

based on work of Drushka et al. (2016), who showed that for the same cumulative rainfall, the maximum rain rate was more important than pulse width in determining the salinity response. Wind speed is applied as a constant value ($U = 0, 2, 4, 8, 12, 16 \text{ m s}^{-1}$) over the four-day time period. The six different rain conditions and six different wind conditions result in a total of 36 model runs. Figure A1 shows an example of idealized forcing and modeled ocean response for one sensitivity run. The model parameters for this study are also listed in Table 1.1.

For event-focused simulations, the initial temperature profile is set as the interpolated profile averaged over five coastal AR events from October 2016–February 2017 from Spray glider data on line 66.7. The initial salinity profile is constant with depth to allow the vertical change in salinity from precipitation to be distinguished from mixing. The salinity at all depths is set to the 10 m salinity from Spray averaged over the same five coastal AR events. The decision to adopt a constant vertical salinity profile is justified by the results of sensitivity tests that indicate that variations in the stratification of the initial vertical salinity profile have little effect on the salinity response to rain events (not shown). In contrast, in a different regime in the tropics, Drushka et al. (2016) and Iyer and Drushka (2021) find that rain falling on saltier water will lead to a larger salinity stratification than rain falling on freshwater, and that the preexisting background salinity can have a larger impact on the salinity response to rain than the rain conditions themselves.

Following Drushka et al. (2016), two metrics are defined in order to characterize the ocean response to rainfall: the depth (D_L) and duration (T_L) of the fresh lens. Here the fresh-lens depth, D_L , is defined as the depth at which the salinity anomaly relative to the salinity at the first time step is 25% of the maximum anomaly. In contrast Drushka et al. (2016) defined D_L where the salinity anomaly relative to a no-rain control run was 10% of the maximum anomaly. The lifetime of the fresh lens, T_L , is defined as the time period over which the fresh-lens depth is non-zero. The definition of D_L differs from that of Drushka et al. (2016) in order to account for AR conditions in the CCS, as ARs in the CCS have smaller rain rates but longer duration than rain events in the tropics. To compare the model simulations for different external forcing cases,

we calculate the salinity difference ΔS as the salinity at 0.01 m depth at each time step subtracted from the 0.01 m depth salinity at the first time step. A positive ΔS therefore represents a decrease in surface salinity over time. The maximum vertical salinity difference, ΔS_{max} , is defined as the maximum value of ΔS within the four-day time period.

Model Setup, Case Studies

Event case studies are run using the one-dimensional configuration of the MITgcm to study the impact of specific AR events on the formation of freshwater lenses. The event length is set to nine days to match the MBARI composite studies. Five different coastal AR events are chosen: (i) 16 October 2016; (ii) 27 November 2016; (iii) 11 December 2016; (iv) 19 January 2017; and (v) 17 February 2017. Atmospheric forcing is applied hourly and is linearly interpolated to 60 s time steps by the model. Values for rain rate (R), wind speed (U_Z and U_M), radiation (I_L and I_S), specific humidity (SpH), and air temperature (T_a) are taken from the ERA5 dataset at the coastal location for a duration starting three days before and ending six days after the event date. Figure A2 shows an example of the forcing for one of the five runs. The initial temperature and salinity profiles are set as the profile for each event starting date from the Spray glider at the coastal location along line 66.7, interpolated to telescoping depths. As in the sensitivity studies, ΔS_{max} is calculated for each model run as the maximum value of the difference in salinity at 0.01 m depth between each time step within the nine-day time period and the first time step. Model output from case studies is compared to that of the sensitivity studies, as well as observational results from the MBARI M1 mooring. The model parameters for this study are also listed in Table 1.1.

Model Validation

A one-dimensional model (the MITgcm ocean column (Adcroft et al., 2018)) will allow for analysis without impacts from horizontal advection or runoff. In order to validate the use of the MITgcm for event-based studies, we first run with external forcing and initial conditions used

by Drushka et al. (2016) for a site in the tropical Pacific and compare with the published results of the General Ocean Turbulence Model (GOTM) by Drushka et al. (2016). For consistency with GOTM outputs, in this model validation ΔS_{max} is defined as the maximum vertical salinity difference between 5 m and 0.01 m, following Drushka et al. (2016). MITgcm results are similar to GOTM results (Fig. A.3). One difference is that the MITgcm KPP tends to mix deeper and preserves the freshwater lens for a shorter duration, except in the case of 10 m s^{-1} winds and 2 mm h^{-1} precipitation rates (not shown). As a result, the maximum vertical salinity difference between 5 m and 0.01 m for a given model run is generally smaller in the MITgcm than in GOTM. Conversely, at higher rain rates, GOTM has greater mixing of large freshwater inputs at the surface, resulting in a lower maximum vertical salinity difference than in MITgcm for 2 m s^{-1} (not shown) winds and 50 mm h^{-1} precipitation rates. However, for most rain and wind cases a statistically significant 1:1 linear fit is exhibited between the two models (Fig. A.3). Therefore differences between GOTM and the MITgcm are judged minor. Since the MITgcm is consistent with the one-dimensional turbulence model, we choose to use it here because it can later be extended to run in a three-dimensional configuration, which will aid in future work considering ocean processes such as horizontal advection, runoff and upwelling.

Sensitivity experiments are run to test other parameters of the MITgcm, including the model time step, the KPP Richardson number threshold for mixing, and the initial stratification (not shown). Model results are relatively insensitive to time step and only sensitive to Richardson number threshold at high rain rates in combination with low wind speeds. Initial stratification is tested by changing the input vertical salinity profile to have different slopes within a salinity range of 33–34 psu in the upper 20–80 m of the water column (not shown). These changes are found to have little impact on the vertical changes in salinity in response to different rain rates.




Figure 1.3. Time series showing salinity (psu) for MITgcm one-dimensional model runs (red, solid) and MBARI M1 Mooring (red, dashed) at 1 m depth, compared to ERA5 rain rate (mm day^{-1}) (blue) from September through March in 2015–2018 (a–d). The black dashed line represents the initial salinity in September for comparison.

1.6 Results

1.6.1 Seasonal Response

While changes in the salinity of the CCS have previously been attributed mainly to advection (Lynn and Simpson, 1987; Schneider et al., 2005), the time series for the MBARI M1 Mooring salinity and the MITgcm model output salinity at 1 m depth in comparison to ERA5 daily cumulative precipitation both suggest that local precipitation also impacts ocean surface salinity (Fig. 1.3). A seasonal freshening is present from September to March for the years 2015–2018 in both mooring and model data, with the exception of 2017 for the mooring (Fig. 1.3). Here, the mooring data often show the freshening to be a response to rain events, as typically spikes in precipitation (10 mm day^{-1} – 35 mm day^{-1}) are followed by decreases in salinity (0.1 psu–1.0 psu). The comparison of model and mooring salinities in Fig. 1.3 shows

that the mooring has a more drastic salinity response immediately following rain events, while the model response is more gradual (up to 0.25 psu). While Fig. 1.3 suggests a relationship between seasonal precipitation and salinity change, its inclusion here is mainly intended as an introduction to the idea that salinity changes in the upper ocean may be linked to precipitation. Data from the MBARI M1 mooring are further analyzed in section 1.6.2.

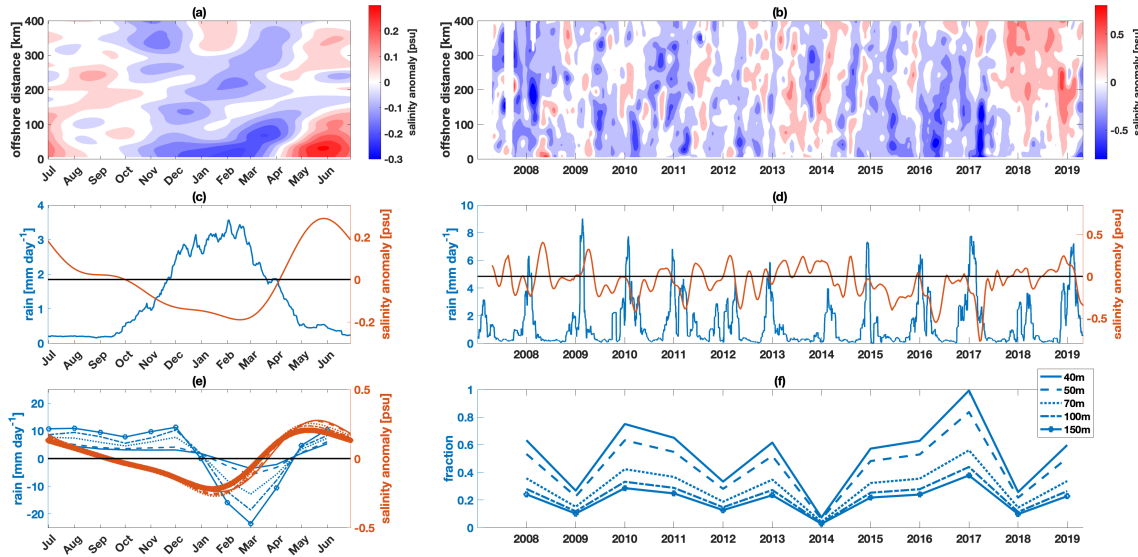


Figure 1.4. (a) Climatological annual cycle and (b) multi-year time series of salinity anomaly as a function of offshore distance at 10 m depth as measured by the CUGN Spray underwater glider on line 66.7. (c,d) Salinity anomaly averaged over offshore distances from 0–50 km (red) and daily precipitation with a 30 day moving mean at the coastal location (blue; offshore distance < 50 km): (c) annual signal averaged over 2007–2019; (d) time series, showing interannual anomaly for salinity and a 30-day moving mean for daily precipitation. (e) Salinity anomaly averaged over different depths (40 m, 50 m, 70 m, 100 m & 150 m) in the upper ocean at 15 km offshore (red) and theoretical daily precipitation that would be required if local rain was the only factor leading to a change in salinity (blue). (f) Ratio of observed cumulative precipitation from September to January of each year to cumulative precipitation that would be required to produce the annual salinity anomaly in (e) for different depths. Spray data from Rudnick et al. (2017a); evaporation and precipitation data from ERA5.

We also examine annual and interannual variability of salinity as measured by the Spray glider and precipitation from ERA5 (Fig. 1.4). The annual climatological salinity anomaly in Fig. 1.4a shows that at all locations there is a negative salinity anomaly (blue) during the

rainy season months of October-April. A positive anomaly (red) is seen during the summer months May-September. This pattern is stronger at the coast than offshore. The annual cycle of negative anomaly in the winter (Oct-Apr) and positive anomaly in summer (May-Sep) is also often visible in the full time series (Fig. 1.4b & d). For example, high precipitation in the 2016–2017 rainy season (Fig. 1.4d) coincides with a negative salinity anomaly (Fig. 1.4d and blue in Fig. 1.4b), while lower precipitation in the 2017–2018 season coincides with a positive, or less negative, salinity anomaly (Fig. 1.4d and red in Fig. 1.4b). Fig. 1.4e shows that the salinity anomaly averaged over the top 40 m to top 150 m is rather insensitive to the depth range over which it is averaged (red lines), suggesting that processes other than local rain (e.g. runoff, advection) play a role in these salinity changes. However, the all-rain scenario is used here as a limiting case by applying these salinity anomalies in Equation (1.1) to calculate the amount of precipitation that would theoretically produce the anomaly if evaporation and rain were the only contributing factors (blue line, Fig. 1.4e). This information is then used to compute the ratio of observed cumulative local precipitation from September to January of each year to the theoretical cumulative precipitation that could account for the annual cycle of freshening. Here, Fig. 1.4f shows that ratio and indicates that local rain could potentially account for up to 100% of the annual cycle of freshening in the upper 50 m in this limiting case in which the system depends only on vertical mixing, with no effect due to horizontal advection. The precipitation required to produce the annual salinity anomaly over the depth range increases with increasing depth, which leads to estimated rain fraction decreasing with increasing integration depth. In other words, as we integrate to greater depth, a smaller portion of the salinity signal is expected to be due to rain. Determining the mechanisms responsible for the residual, which possibly include horizontal advection, runoff, upwelling, or downwelling, is outside the scope of this study.

To characterize upper-ocean freshening in response to precipitation, for both glider and model data, we plot the March-minus-September salinity differences at 10-m depth as a function of cumulative rainfall at coastal, onshore, and offshore locations (Fig. 1.5 a–c). We also include salinity differences as measured from the MBARI M1 mooring at the coastal location. The

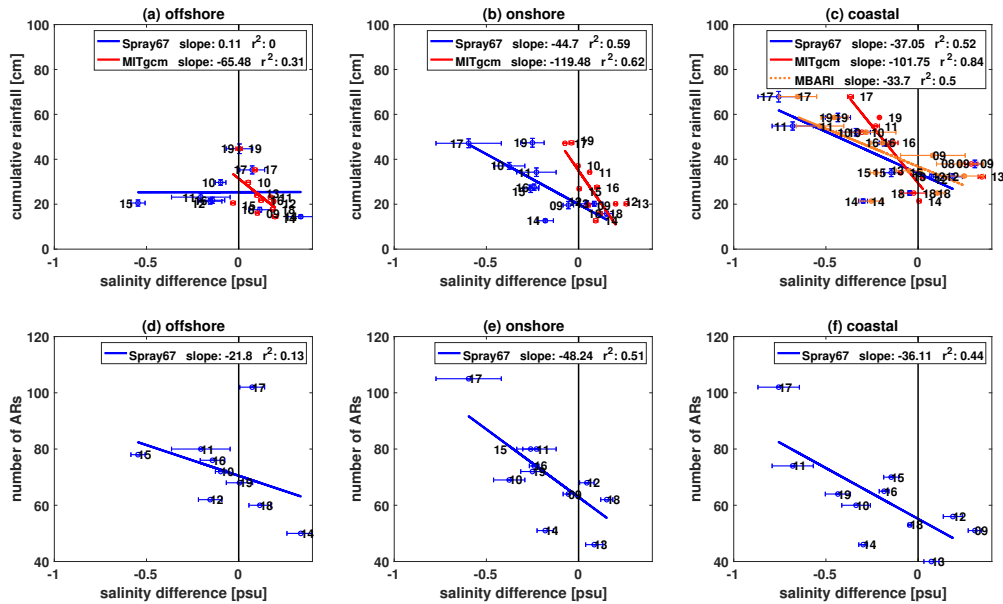


Figure 1.5. (a-c) Cumulative rainfall (cm) and (d-f) number of AR events as a function of salinity (psu) difference between March and September for the years 2008–2019 at (a,d) offshore, (b,e) onshore, and (c,f) coastal locations. Panels a–c include CUGN Spray line 66.7 observations (blue), MBARI M1 mooring observations (red, dotted) and MITgcm one dimensional model runs (red, solid) at 10 m depth. The blue and red lines represent least squares fits to glider, mooring, and model data with the slope and r^2 values labeled in the legend. Panels d-f show data from SIO-R1 AR catalog and CUGN Spray line 66.7 observations (blue) at 10 m depth. Blue lines represent linear regressions, with slopes and r^2 indicated in the legends.

quantities appear anti-correlated: high cumulative rainfall typically corresponds to larger salinity decreases (Fig. 1.5a–c). For glider, mooring, and model data, least squares fits show negative slopes and r^2 values that are statistically significant at the 95% level (corresponding to $r^2 > 0.30$ for 12 years of data), except at the offshore location. These r^2 values suggest that precipitation can explain a significant portion of the variance in salinity difference over the rainy season at coastal and onshore locations (52% and 59% for the glider data, 50% for the mooring data, and 84% and 62% for offshore). The offshore region does not always show a salinity decrease over the course of the water year, and it also tends to experience a lower cumulative rainfall than coastal and onshore locations (15–45 cm for offshore in comparison to 20–70 cm for coastal). The model response differs from the observational data in that the model tends to show

a smaller decrease in salinity over the season (Fig. 1.5a–c), as discussed in section 1.5.1.

Given the one-dimensional nature of the model, external forcing would be expected to explain 100% of the variance in salinity changes, which is not the case in Fig. 5. Here, unexplained variance results from not including evaporation and analyzing salinity changes only at the surface, thus not capturing mixing of the freshwater input to further depths. When comparing evaporation minus precipitation to the salinity change integrated over all depths, 100% of the variance is explained by the model for all locations (not shown).

To further investigate the role that ARs play in seasonal upper-ocean freshening, we compare the number of AR events to the March-minus-September 10-m salinity difference for glider data at the three locations (Fig. 1.5d–f). Years with more ARs tend to exhibit larger salinity decreases, as seen in Fig. 1.5d–f and as indicated by the negative slopes of the regressions. This is the case except in 2017, when an increase in salinity is seen despite a large number of ARs (Fig. 1.5d). Similarly to the relationship between cumulative rainfall and salinity difference, this trend is statistically significant at the 95% level, except at offshore locations, and r^2 values suggest that ARs can explain a significant portion of the variance in salinity difference over the rainy season for coastal and onshore locations. At offshore locations, relationships between the number of AR events and salinity difference (Fig. 1.5d) or precipitation and salinity difference (Fig. 1.5a) do not exhibit r^2 values for linear regression that are statistically significant. The lack of correlation between local rainfall and freshening at offshore locations could be caused by salinity changes related to processes other than rainfall, such as advection.

1.6.2 Event-Based Response

Event Composites

While the results of section 1.6.1 demonstrate that in the CCS region, the upper ocean freshens more during high rainfall years than it does in low rainfall years, the question of whether individual rainfall events are detectable in upper-ocean salinity remains. We begin examination of the ocean salinity response to rain events on short time scales by using event composites.




Figure 1.6. Composite time series of (a) six-hourly rain (mm, red) and daily cumulative rain from day -3 to day n (mm, blue), (b) wind speed (m s^{-1}) with a six-hour moving mean, and (c) salinity difference (psu) between relative day n and relative day 0 for 85 rain events occurring at the MBARI M1 mooring location from January 2007 – March 2019. The solid line (μ) represents the mean of all composite events and the shading represents the standard error of the mean (sem) among these events. The solid black line in (c) represents a salinity of zero, which is zero on day zero because the anomaly is in reference to this day. Events are included if daily cumulative precipitation on day zero is greater than 5 mm day^{-1} and there has not been another rain event within 10 days of the event start date. Event start dates are set as the first date that rainfall exceeds the threshold; conditions are shown from 3 days before through 6 days after this date. Rainfall and wind speed are taken from ERA5 and salinity from the MBARI M1 Mooring.

Figure 1.6 shows a time series composited from 85 events that occurred at the MBARI M1 mooring location from January 2007–March 2019 (see Fig. A.4). The rain events that are in the composite analysis are shown as both cumulative rain over six hours (red) and daily cumulative precipitation (blue), whereas salinity is plotted as a six-hourly moving mean. In Fig. 1.6, relative day zero represents the first day that rainfall exceeded a threshold of 5 mm day^{-1} (a result of the event compositing discussed in section 1.5.2). The wind speed (Fig. 1.6b) remains relatively constant at about $5 \pm 1 \text{ m s}^{-1}$ for the duration of the composite time series, with a slight peak on relative days 0–1. Figure 1.6c shows that the surface salinity measured by the M1 mooring decreases over the duration of the composite time series, especially during the days with peak rain (day 0 through 1). While there is an increase in salinity from day 1 through day 4, overall the salinity is lower at the end of the composite time series than at the beginning. The results

from this composite study indicate that salinity measurably decreases in response to rain on an event basis. To assess the mechanisms governing this freshening pattern, we use the model to carry out event sensitivity studies.

Model Sensitivity Studies of Rain and Wind Effects in Freshwater Lens Formation

Event-based studies are performed using the one-dimensional MITgcm configured for the CCS. The model allows us to isolate the impacts of rain and wind on upper-ocean salinity stratification and to determine whether the resulting vertical salinity change will be detectable, given the 0.01 psu resolution of CTD instruments (as discussed in section 1.4.1). While the range of salinity responses depends on rain rate and wind speed on event time scales, this study highlights two key mechanisms that govern salinity changes as a function of precipitation and wind speed: (i) mixing of the freshwater or (ii) development of freshwater lenses at the surface.

Figure 1.7 shows the salinity anomaly in the upper ocean in response to a range of model input conditions (wind speeds increase from 2 to 16 m s⁻¹ from top to bottom, and rain rates increase from 2 to 8 mm h⁻¹ from left to right), normalized to the maximum salinity anomaly for each given wind speed and rain rate. Two extreme cases are detected: (i) vertical mixing of the freshwater to depths greater than 20 m at high wind speeds ($U > 8$ m s⁻¹) and (ii) development of freshwater lenses at the surface for low wind speeds ($U \leq 8$ m s⁻¹), where the depth of the fresh lens is depicted by the black lines of Fig. 1.7. This is consistent with results from Thompson et al. (2019), where stable rain layers were found to persist with wind speeds up to 9.8 m s⁻¹. As wind speed increases (moving top to bottom) the freshwater lens is brought to a greater depth and remains over a shorter time period than at low wind speeds, except in the case of $R = 2$ mm h⁻¹ where the small freshwater input may impact the trend in lens depth. As rain rate increases (moving left to right) the freshwater input is mixed over a deeper range, except in the case of $U = 2$ m s⁻¹; additionally the lens has a longer duration. These results are reproduced in Fig. 1.8.

The dependence of the vertical salinity gradient on rain and wind speed is shown in

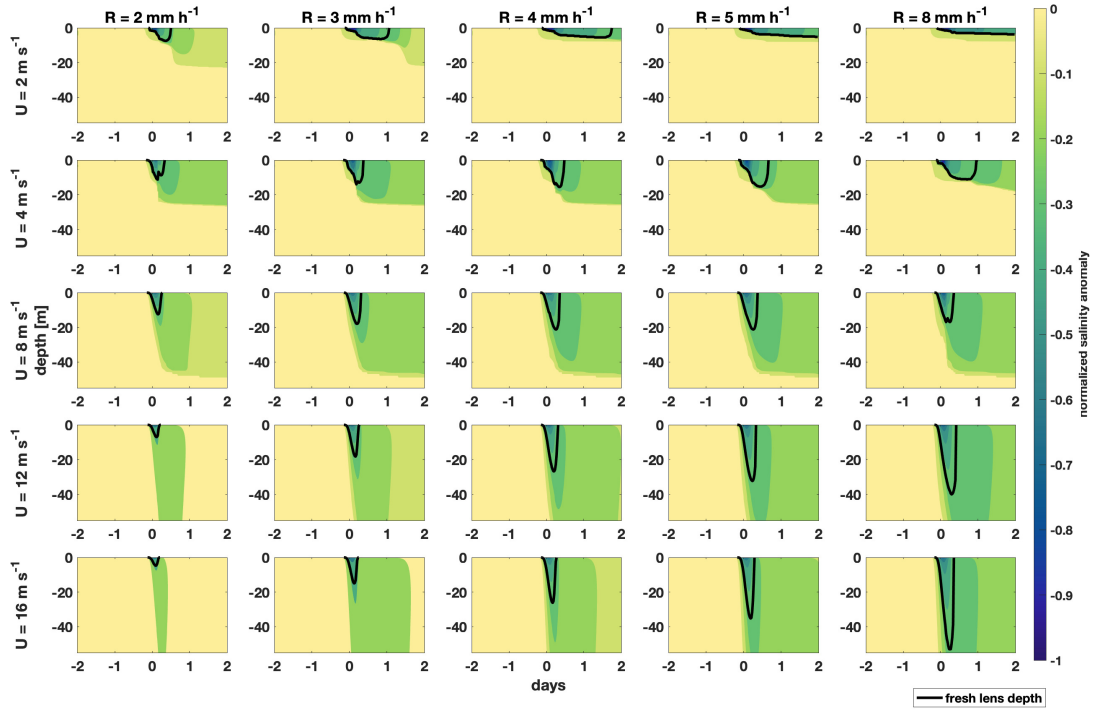


Figure 1.7. Normalized salinity anomaly in the upper 55 m of the ocean for the four-day one-dimensional MITgcm runs and for wind speeds from 2 to 16 m s^{-1} and maximum rain rates from 2 to 8 mm h^{-1} . Each contour plot is divided by the absolute value of the maximum salinity anomaly for the given rain rate and wind speed. Black lines represent the freshwater lens depth, D_L (m), defined as the depth at which the salinity anomaly relative to the salinity during the first time step for each run is 25% of the maximum anomaly.

Fig. 1.8. In Fig. 1.8a & b, the maximum vertical salinity difference, ΔS_{max} (defined in section 1.5.2), increases as a function of rain rate and decreases as a function of wind speed. Modeled freshwater lens depth (D_L) and duration (T_L) are shown as a function of wind speed and rain rate in Fig. 1.8c-d. Here, an increased wind speed corresponds to deeper mixing, bringing freshwater to a greater depth, therefore decreasing stratification and decreasing the magnitude of ΔS_{max} . At low wind speeds there is minimal mixing, and changes in salinity are confined to the surface (<20 m) and are not prominent at depth, leading to a larger ΔS_{max} (Fig. 1.8a & b). In this case, a freshwater lens is formed at the surface, and stratification is enhanced. Figures 1.8a & b (reproduced in Fig. 1.9) also show model output from five event case studies (the colored circles),

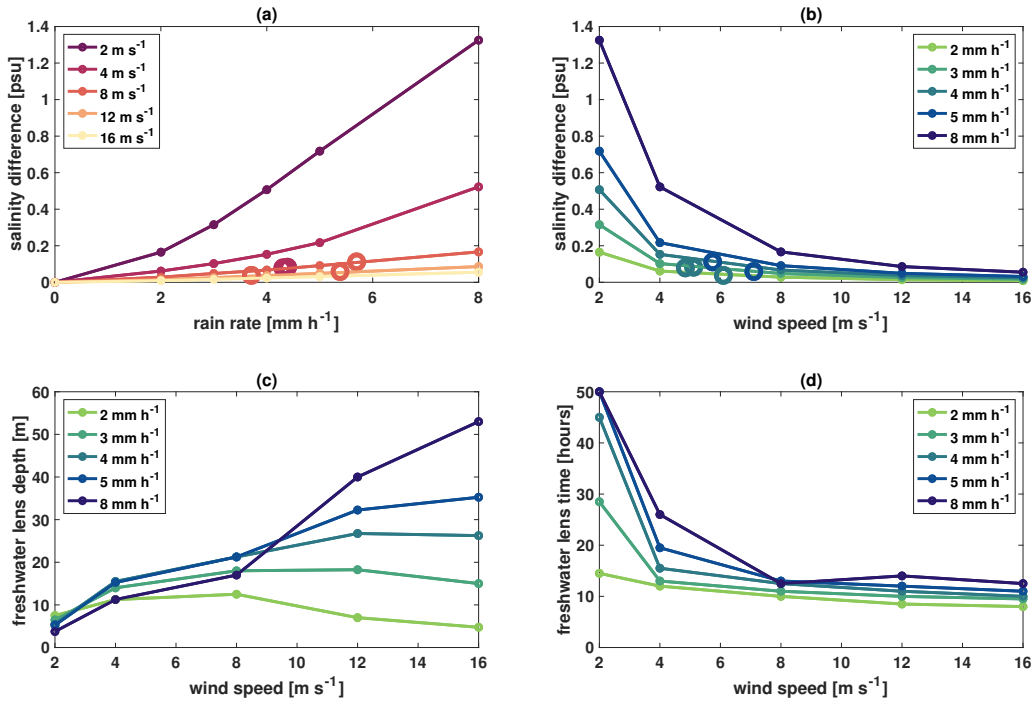


Figure 1.8. Results from the MITgcm experiments using idealized environmental forcing in which the peak rain rate and the wind speed are varied. (a) Peak magnitude of ΔS , ΔS_{max} , as a function of rain rate for five different wind speeds; (b) ΔS_{max} as a function of wind speed for different rain rates; (c & d) maximum (c) thickness, D_L , and (d) lifetime, T_L , of the fresh lens as a function of wind speed at different rain rates. ΔS_{max} is defined as the maximum value of the salinity difference at 0.01 m depth from the salinity at the first time step within the four-day simulation time period. In both figures (a) and (b), the colored circles show model output from event case studies, with the colors representing wind speed and rain rate, respectively.

which fall within the same range for ΔS_{max} as the output from the sensitivity studies with similar rain rates and wind speeds. The black dotted line in Fig. 1.9 represents the salinity change that is detectable by CTD instruments (0.01 psu). Almost all of the events in the sensitivity studies exceed this threshold, with the only exception being for a rain rate of 2 mm h^{-1} in combination with a wind speed of 16 m s^{-1} .

The results show a relationship between wind, rainfall, and salinity similar to that suggested by Drushka et al. (2016): $\Delta S_{max} = AR_{max}U^b$, where constants A and b are solved for using model outputs. Here, rain rates of 0 mm h^{-1} and wind speeds of 0 m s^{-1} are omitted from the regression because the fit is representative of cases where rain and wind are present. For the

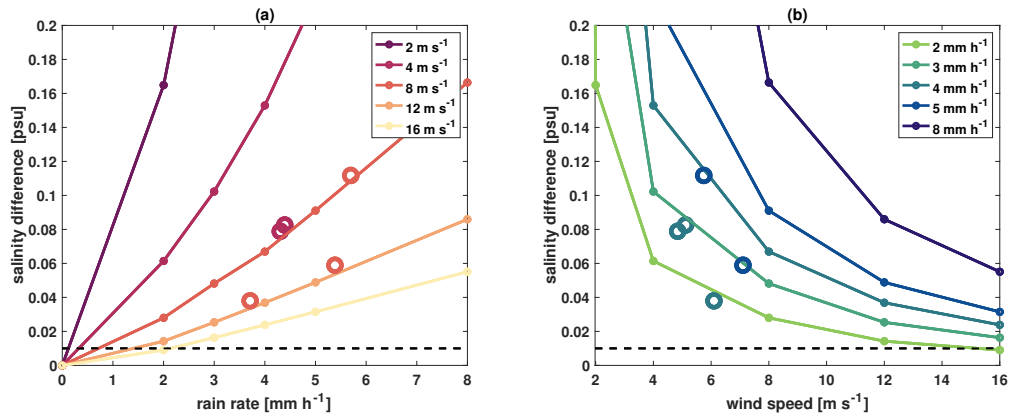


Figure 1.9. Same as Fig. 1.8 a & b, zoomed in to enhance view of results from event case studies (colored circles). The colored circles show model output from event case studies, with the colors representing wind speed and rain rate, respectively. The black dotted line represents the salinity difference of 0.01 psu that is detectable by CTD instruments.

MITgcm model runs, $A = 0.32 \pm 0.05 \text{ psu (mm h}^{-1}\text{)}^{-1}$ and $b = 1.44 \pm 0.06$. Uncertainties of linear regression parameters are calculated using Monte Carlo methods (Fig. A.5). The values of the regression parameters are within five standard deviations of values found by Drushka et al. (2016): $A = 0.11 \pm 0.03$ and $b = 1.1 \pm 0.03$. The values of these coefficients are also similarly related to those found in studies done without the wind dependence both by Drucker and Riser (2014), who found a value $A = 0.14 \text{ psu (mm h}^{-1}\text{)}^{-1}$ averaged over the tropics, and by Boutin et al. (2014), who found region-dependent values of A that ranged from 0.14 to 0.22 $\text{psu (mm h}^{-1}\text{)}^{-1}$ at moderate wind speeds. Differences in these coefficients likely arise as a result of the difference in duration of the applied rain pulse (12 h here for AR studies in CCS versus 1 h for studies in the tropics). While this relationship has been applied in the tropics for the references listed above, we find it does well in representing AR events in the CCS, with an r^2 of 0.97 (Fig. A.5). It should be noted that this equation is appropriate for one-dimensional models that do not include advection, and may not work well in cases where advection is significant. However, case studies in the following section (section 1.6.2) show this equation does well in representing the magnitude of the salinity response to AR events in comparison to in situ measurements (Figs. 1.9 & 1.10).

Freshwater lenses reach depths of 5–50 m, depending on rain rate and wind speed (Fig. 1.8c). The depth of the fresh lens increases with wind speed for all rain rates, except in the cases of 2 mm h^{-1} and 3 mm h^{-1} rain rates where wind is greater than 8 m s^{-1} . These exceptions likely occur because the freshwater input is too small to cause salinity changes at increasing depths during mixing. Additionally, the fresh lens depth increases with higher rain rates, as indicated by the ordering of the green lines, with the lowest rain rate (light green, 2 mm h^{-1}) having the smallest D_L and the highest rain rate (dark blue, 8 mm h^{-1}) the largest D_L . This is true except in the cases of low wind speed and high rain rate ($U = 2, 4 \text{ \& } 8 \text{ m s}^{-1}$ and $R = 8 \text{ mm h}^{-1}$), where the magnitude of the salinity response is comparatively large ($\Delta S_{max} = 1.3, 0.55 \text{ \& } 0.2 \text{ psu}$). These events fall outside the trend for D_L because for each particular combination of wind speed and rain rate these metrics are defined based on the maximum salinity anomaly relative to the salinity at the first time step, which for these extreme cases is much higher than the average salinity anomaly for a particular rain rate or wind speed. Freshwater lenses last anywhere from 10–50 h, depending on rain rate and wind speed (Fig. 1.8d). The duration of the freshwater lens, T_L , shows a pattern of decreasing with increasing wind speed and decreasing rain rate. For wind speeds greater than 8 m s^{-1} the lens duration has a much smaller range of 10–15 h.

Results for the fresh lens depth, D_L , are in agreement with the the 20 m mean stable layer depth in central Indian Ocean found by Thompson et al. (2019). These results also show similar trends to the tropical results of Drushka et al. (2016). One difference is that for these studies of characteristic AR events in the CCS, the depth and duration of the freshwater lens are much larger than studies done in the tropics. This is likely a result of the fact that AR events in the CCS have a much longer rainfall duration than rain events in the tropics (12 h versus 1 h). This is confirmed by runs done in the CCS with 24 h rain pulses (not shown), where D_L and T_L increased even more from the 12 h rain pulse case. It should be noted that D_L and T_L are highly sensitive to the lens definition, as discussed in section 1.5.2. Decreasing the percentage of the maximum salinity anomaly that defines the depth leads to overall increases in D_L and T_L . This makes sense because a less drastic salinity anomaly is expected to reach greater depths for a

longer duration. As an example of this sensitivity, for a rain rate of 8 mm h^{-1} and $U = 12 \text{ m s}^{-1}$, when D_L is defined as the depth at which the salinity anomaly is 15% of the maximum anomaly, rather than 25%, it reaches a maximum of 80 m instead of 53 m. Correspondingly, the time, T_L , reaches a maximum of 95 h instead of 50 h.

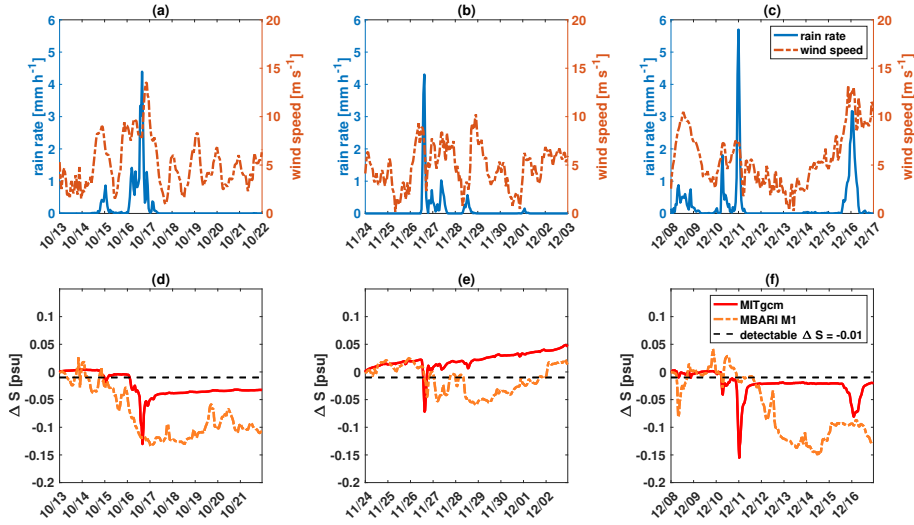


Figure 1.10. Results from case studies for three AR events in the CCS. (a–c) Time series of rain rate (mm h^{-1} , blue) and wind speed (m s^{-1} , red) that was used as model forcing from ERA5 at the coastal location; (d–f) time series showing the salinity difference (ΔS , psu) from the first time step at 0.01 m depth for the model output (red, solid) and at 1 m depth for the MBARI M1 mooring (orange, dotted). The black dotted line in (d–f) indicates the salinity difference of 0.01 psu that is detectable by CTD instruments. The start date for event 1 (a,c) is 16-OCT-2016; event 2 (b,e) is 27-NOV-2016; and event 3 (c,f) is 11-DEC-2016. The model runs were initialized three days before this date, and run until six days after.

Model Case Studies

Event case studies are performed using the one-dimensional MITgcm configured for the CCS at the start of each of five different AR events (Table 1.1). The model allows us to isolate the impacts of atmospheric forcing on upper-ocean salinity stratification and to determine whether the resulting vertical salinity change may be detectable, given the 0.01 psu resolution of CTD instruments (as discussed in section 1.4.1). The results from three case studies are shown in Fig. 1.10, where the different columns (i.e. a&d, b&e and c&f) represent each of the three

different events. The top row (a–c) shows the rain rate (blue) and wind speed (red) from ERA5 at the coastal location that was used as forcing for the model. The second row (d–f) shows the response of salinity difference (ΔS) from the first time step at 0.01 m depth for the model (red, solid) and 1 m depth for the MBARI M1 mooring (orange, dotted). The magnitude of the model and mooring ΔS responses are similar, while their temporal structure is not. The mooring often has a slower response that lasts a longer duration. These differences are likely due to the fact that the model is one-dimensional and solely shows a salinity response to rain, while the mooring captures runoff and advection of waters from other locations that were impacted by the rain events, and thus changes continue to occur once the local rain has stopped. Here, the black dotted line indicates ΔS values that are detectable by CTD instruments (0.01 psu), showing that all three AR events produced measurable changes in salinity. Additionally, Fig. 1.9 shows the results from five modeled case studies overlaid on results from the model sensitivity studies (colored circles), as a function of both rain rate and wind speed. The black dotted line indicates ΔS values that are detectable by CTD instruments (0.01 psu). All of the the case studies shown produce salinity changes greater than the measurable threshold. The ΔS values for the case studies fall within the range of the sensitivity studies for a given rain rate and wind speed, as discussed in section .1.6.2. Overall, the salinity difference, ΔS , in the modeled case studies is consistent with outputs from the model sensitivity studies for characteristic AR events, as well as with observations at the MBARI M1 mooring.

1.7 Discussion

The purpose of this study has been to evaluate the impact of atmospheric forcing on surface ocean salinity in the CCS. A one-dimensional ocean model can help isolate the salinity response to rainfall events in comparison to other intrinsic ocean dynamics. While changes in salinity in the CCS have previously been largely attributed to southward horizontal advection of low salinity water from the northeast Pacific (Lynn and Simpson, 1987; Schneider et al., 2005),

this analysis has shown that the salinity changes could also be attributed to freshwater inputs in the form of precipitation from atmospheric rivers on both seasonal and event timescales.

1.7.1 Seasonal Response

Seasonal freshening in the CCS depends on cumulative rainfall. Results in section 1.6.1 compare ERA5 rainfall to salinity from observational data (mooring and underwater glider) and one-dimensional model output. While intrinsic ocean processes should be captured by observations, most are not represented by the one-dimensional model. Despite this omission, the model nonetheless shows a statistical relationship between cumulative rainfall and salinity difference (Fig. 1.5). These analyses support the idea that local rainfall may be one of several mechanisms playing a role in the seasonal salinity response, and that it is a significant enough component to account for anomalously fresh or salty years.

We find that there is a stronger salinity signal in coastal locations for both observations and model outputs. As discussed in section 1.6.1, this could be attributed to the fact that there is a higher cumulative rainfall at coastal locations. Additionally, processes omitted by the model, including upwelling, runoff and advection, could all play a role in the observational results. For example, Auad et al. (2011) suggest that upwelling of cool, saline water enhances coastal salinity increases in the summer, which could contribute to a larger positive salinity anomaly in summer (September) and a larger difference in March minus September salinity. Freshwater input from riverine runoff has also been linked to decreases in surface salinity measurements. AR precipitation events occur more often on land than over the ocean (Fig. 1.1a), which might lead to runoff. Riverine input from the Salinas River that discharges into Monterey Bay has been linked to decreases in surface salinity as measured by the MBARI M1 Mooring (Kudela and Chavez, 2004). River discharge from the Sacramento/San Joaquin River system 100 km north of the M1 Mooring has also been linked to low salinity measurements off the coast of Monterey Bay (Johnson et al., 1999).

Southward advection of freshwater in the low-salinity tongue of the California Current

has been previously described as the main source of salinity changes in the CCS (Aquad et al., 2011; Lynn and Simpson, 1987; Schneider et al., 2005). While we do not find evidence against this, when looking at the seasonal cycle of CCS advection there are a few instances of anomalous salinity that may not be linked to advection. For example, the low surface salinity anomaly seen 50 m offshore along CalCOFI line 66.7 during the winter months (Fig. 4.2.3.1 in Rudnick et al., 2017b) is unexplained by the strong poleward current at this location and time which would be expected to carry saltier water from further south. On longer timescales (5–10 years), Schneider et al. (2005) found that negative anomalies in salinity storage averaged over the top 150 m corresponded to increased precipitation, but also noted that patterns in salinity anomaly imply freshwater fluxes that are larger than the observed precipitation or evaporation anomalies. This is supported by Fig. 4f, which shows that the observed precipitation is 3–30% of the precipitation that would be required to produce the salinity anomaly in the upper 150 m if all other terms in the salinity balance are ignored. While this may be the case for the salinity changes in the upper 150 m, we have shown the observed precipitation can explain up to 100% of the seasonal salinity change in the upper 40 m.

While some of the salinity changes may be linked to runoff, upwelling, or advection, the one-dimensional nature of the model omits these ocean dynamics that might have a visible impact on mooring and glider data. Nonetheless, the model still shows a seasonal salinity response to freshwater inputs from rain, as discussed in section 1.6.1.

1.7.2 Event-Based Response

On event time scales, certain combinations of rain rate and wind speed can lead to the formation of freshwater lenses. Freshwater lenses may inhibit mixing of surface waters and increase upper-ocean stratification, which has a variety of implications for the exchange of heat and moisture between the ocean and atmosphere, as discussed in section 1.3.4 (SPURS-2 Planning Group, 2015; Williams et al., 2006). Understanding the structure and evolution of these lenses is important for understanding the possible impacts on air–sea exchanges.

The wind speed and rain rate dependences of ocean surface salinity are investigated using event composites and one-dimensional model sensitivity studies. We show that salinity decreases in response to rain events (section 1.6.2). Furthermore, model results show that the salinity change during a rain event depends linearly on the rain rate, and is inversely proportional to wind speed (section 1.6.2). This suggests that for low wind speeds, freshwater inputs are trapped at the surface and lead to the formation of freshwater lenses, while high wind speeds cause freshwater from rain to mix as deep as 50 m and prevent the formation of long-lasting fresh lenses.

Many events characteristic of ARs in the CCS produce measurable changes in salinity. As discussed in section 1.6.2, there is only one instance where the sensitivity studies do not produce a salinity change that exceeds the 0.01 psu detectable limit (low rain rate in combination with high wind speed). Additionally, all modeled and observed case studies produce measurable salinity changes. Case studies show that single AR events can produce salinity decreases of up to 0.1 psu that last up to 50 hours (Fig. 8). These salinity anomalies are comparable to the decreases in salinity over the entire rainy season, which are shown to be as high as 0.8 psu for observations, and 0.4 psu for one-dimensional models where effects from advection, runoff and upwelling are excluded (Fig. 5). It should be noted that while a single AR event may not cause a large, long-lasting drop in salinity, there is a range of salinity change depending on the strength of the given AR. Additionally, ARs often occur in series with several in a row, which may lead to a larger integrated effect over time. Statistics from a composite analysis of 91 AR events from Table 2 of Ralph et al. (2013) indicate that the average maximum rain rate for these events is 4.09 mm h^{-1} and the average wind speed is 12.8 m s^{-1} . Based on our results, these events would produce salinity changes above the measurable threshold, implying that AR events should be detectable by CTD measurements of ocean salinity.

1.8 Conclusion

Seasonal freshening in the CCS depends on cumulative rainfall and atmospheric river events, in addition to other intrinsic ocean dynamics that previous studies have identified. At coastal and onshore locations, the CCS freshens throughout the rainy season due to AR events, and years with higher AR activity are associated with a stronger freshening signal (Fig. 1.5).

Event studies indicate that freshening in the CCS depends on wind speed in addition to rain rate. Low winds lead to conditions that cause freshwater lens formation, while high wind speeds mix freshwater input from rain through the mixed layer. Results from our one-dimensional model show that freshwater lens formation in the CCS is possible in the event of heavy rain and low winds. For events that are characteristic of ARs in the CCS, these lenses are formed often and can last anywhere from 10–50 h. The one-dimensional model simulations also suggest that events characteristic of ARs in the CCS tend to produce changes in salinity that are greater than the measurable CTD limit of 0.01 psu, as indicated in Figs. 1.9 & 1.10.

Because of the dependence of salinity on both rain and wind, further investigation in the CCS would require local, high-resolution observations of both variables, as was done in the SPURS-2 experiment, in order to develop a more complete understanding. With observations it would also be possible to validate the use of the one-dimensional MITgcm to represent salinity changes on an event time scale, as was done for the seasonal studies (e.g. Fig. 1.2 in section 1.5.1).

As discussed in section 1.6.2, the freshwater lens is highly sensitive to definition. The definitions for D_L and T_L that were shown to work with GOTM for the salinity response to rain events in the tropics (Drushka et al., 2016) were altered slightly for results in the CCS, as discussed in section 1.5.2. In another study, Thompson et al. (2019) derived an estimate of the stable layer depth based on wind speed and buoyancy frequency. Future work could explore different forms of the definition specific to the CCS.

While this study has provided evidence that freshwater inputs from rain contribute to

variability in ocean surface salinity, the relative importance of horizontal advection, runoff, and external atmospheric forcing has not been addressed. Advection could contribute to the evolution of freshwater lenses by causing increased mixing and by introducing new water into the region. Future studies could address these shortcomings by considering a three-dimensional ocean model that will show the relative importance of horizontal advection and runoff. Additionally, large-scale surface advective salinity transport could be estimated from observations. Future work could also look at the response of properties other than salinity, for example temperature or biogeochemical properties, and thus elucidate the impact of precipitation events on the climate state.

1.9 Acknowledgments

Figures in this report were prepared using MATLAB, Matplotlib: A 2DGraphics Environment Hunter (2007). DG was supported by NSF Award 1928305. MRM and STG were supported by NASA awards NNX16AH67G and 80NSSC20K1136. DG and STG were supported by NASA award 80NSSC19K0059. Special thanks to the coauthors for their invaluable contributions to this chapter: Matthew Mazloff, Sarah Gille, Donata Giglio, and Aniruddh Varadarajan. We thank our reviewers for their helpful feedback.

Chapter 1, in full, is a reprint of the material as it appears Journal of Physical Oceanography 2023. This chapter is published as: Hoffman, L., Mazloff, M. R., Gille, S.T., Giglio, D., and Varadarajan, A. (2022). Ocean Surface Salinity Response to Atmospheric River Precipitation in the California Current System. *Journal of Physical Oceanography*, 52(8): 1867–1885. ©American Meteorological Society. Used with permission. The dissertation author was the primary investigator and author of this paper.

1.10 Data Statement

We acknowledge all sources of publicly available data that were used in this study. This paper contains modified Copernicus Climate Change Service information [2020] in the form of ERA5. The European Commission nor ECMWF is responsible for any use that may be made of the Copernicus information or data it contains. The ERA5 dataset can be accessed at <http://doi.org/10.24381/cds.e2161bac>. Data from the CUGN were produced by Daniel Rudnick at Scripps Institution of Oceanography and can be accessed at <https://spraydata.ucsd.edu/projects/CUGN/>. Mooring data were made available by the Monterey Bay Aquarium Research Institute (MBARI) and can be accessed at <https://www.ncei.noaa.gov/archive/accession/0130040>. The MITgcm one-dimensional model was made available by the MITgcm contributors Adcroft et al. (2018). The SIO-R1 Atmospheric River Catalog can be accessed at <https://weclima.ucsd.edu/data-products/>. All of the data and code used for processing for this paper can be accessed at <https://doi.org/10.6075/J0BV7GGW>.

Table 1.1. Model parameters for (a) seasonal (b) event sensitivity and (c) event case studies.

Study Time scale	(a) Seasonal	(b) Event Sensitivity	(c) Case Studies
Model Parameters (one-dimensional MITgcm)			
Time step (seconds)	1800	60	60
Run time (days) / number of time steps	213 / 10244	4 / 5760	9 / 13020
Depth (m) / dZ (m)	140 / 0.5	140 / telescoping	140 / telescoping
External forcing input interval (seconds)	86400	60	3600
Number of runs	13 (September–March, 2008–2019)	36 (six rain rates / six wind speeds)	five (16 October 2016, 27 November 2016, 11 December 2016, 19 January 2017, 17 February 2017)
Initial Conditions (from Spray)			
Salinity profile	averaged over September for each year within each off-shore distance regime (coastal, onshore, offshore)	constant from salinity average over five coastal AR events at 10 m depth, telescoping depths	salinity on event start date at coastal location, interpolated to telescoping depths
Temperature profile	averaged over September for each year within each off-shore distance regime (coastal, onshore, offshore)	temperature average over five coastal AR events, interpolated to telescoping depths	temperature on event start date at coastal location, interpolated to telescoping depths
External Forcing (from ERA5)			
Rain rate	daily cumulative	idealized 12 h Gaussian pulse (0, 2, 3, 4, 5, & 8 mm h ⁻¹)	hourly
Wind speed	daily mean	idealized constant over four days (0, 2, 4, 8, 12, & 16 m s ⁻¹)	hourly
Atmospheric temperature, specific humidity, short and longwave radiation	daily mean	constant (T_a , 13.1°C ; SpH , 0.008 kg kg ⁻¹ ; I_s , -106.3 W m ⁻² ; I_L , -323.2 W m ⁻²), average over five AR events at the coastal location	hourly

Chapter 2

Machine learning for daily forecasts of Arctic sea-ice motion: an attribution assessment of model predictive skill.

2.1 Summary

Physics-based simulations of Arctic sea ice are highly complex, involving transport between different phases, length scales, and time scales. As a result, numerical simulations of sea-ice dynamics have a high computational cost and model uncertainty. We employ data-driven machine learning (ML) to make predictions of sea-ice motion. The ML models are built to predict present-day sea-ice velocity given present-day wind velocity and previous-day sea-ice concentration and velocity. Models are trained using reanalysis winds and satellite-derived sea-ice properties. We compare the predictions of three different models: persistence (PS), linear regression (LR), and convolutional neural network (CNN). We quantify the spatio-temporal variability of the correlation between observations and the statistical model predictions. Additionally, we analyze model performance in comparison to variability in properties related to ice motion (wind velocity, ice velocity, ice concentration, distance from coast, bathymetric depth) to understand the processes related to decreases in model performance. Results indicate that a CNN makes skillful predictions of daily sea-ice velocity with a correlation up to 0.81 between predicted and observed sea-ice velocity, while the LR and PS implementations exhibit

correlations of 0.78 and 0.69, respectively. The correlation varies spatially and seasonally; lower values occur in shallow coastal regions and during times of minimum sea-ice extent. LR parameter analysis indicates that wind velocity plays the largest role in predicting sea-ice velocity on one-day time scales, particularly in the central Arctic. Regions where wind velocity has the largest LR parameter are regions where the CNN has higher predictive skill than the LR. ¹

2.2 Introduction

Sea-ice cover in the Arctic has been diminishing since the beginning of the satellite record (Serreze et al., 2007; Stroeve et al., 2012; Stroeve and Notz, 2018; Thoman et al., 2022). Negative trends in sea-ice concentration, thickness, and multiyear ice coverage (Carmack et al., 2015) have been reported throughout the Arctic, whereas the length of the melt season, drift speeds, and deformation rates are increasing (Stroeve and Notz, 2018; Rampal et al., 2009; Onarheim et al., 2018). Climate model simulations indicate a substantial likelihood that the Arctic Ocean will become largely ice free during September by 2100 if warming exceeds 2°C (Stroeve and Notz, 2018; Notz and Stroeve, 2018; Jahn, 2018; Meredith et al., 2019). Transition to thinner and more fragile ice will have widespread environmental, geopolitical, and logistical impacts, including potential for new increased maritime activity (Bennett et al., 2020; Crawford et al., 2021; Cao et al., 2022), with which comes the need to know where sea-ice is and the need for skillful predictions of where it will be. In this study we contribute to addressing these issues by assessing the skill of machine learning models in making one-day predictions of sea-ice motion. We design these models to predict present-day ice motion based on previous-day observations, and show proof-of-concept for applications in operational forecasting that would allow information about the ice state to be obtained before satellite retrievals are processed. Additionally, we explore the extent to which these ML models will have enough skill to be used to represent the dynamical

¹This chapter is published as Hoffman, L., Mazloff, M. R., Gille, S.T., Giglio, D., Bitz, C. M., Heimbach, P., and Matsuyoshi, K. (2022). Machine learning for daily forecasts of arctic sea-ice motion: an attribution assessment of model predictive skill. *Artificial Intelligence for Earth Systems, in review*. ©American Meteorological Society. Used with permission.

component of sea ice in a simulation framework that provides nowcasting of the state of Arctic sea ice.

Predictions of sea-ice motion have almost exclusively been attempted with numerical prediction models (Petrou and Tian, 2019). While these state-of-the-art physics-based models for sea-ice prove useful, their inherent complexity comes with a high computational cost (Hunke et al., 2020). There are also several sources of uncertainty, including large sensitivity to initial conditions and physical assumptions (Blanchard-Wrigglesworth et al., 2015). In contrast to physics-based models, machine learning is emerging as a powerful tool for applications in the geosciences in cases where large volumes of data are available (Hsieh and Tang, 1998; Toms et al., 2020). Machine learning predictions are driven by data and therefore do not depend on assumptions imposed on physical constraints. Although these constraints are crucial for some applications (e.g. where mass, heat, and momentum need to be conserved), in other applications they introduce additional uncertainty and complexity with little scientific benefit. While simple forms of machine learning (e.g. linear regression) have been commonly used in the geosciences, more advanced deep learning models (e.g. neural networks) have the potential to further elucidate physically meaningful relationships within data (McGovern et al., 2019; Toms et al., 2020). In this study, we assess the viability of using a neural network as a surrogate model to parameterize sea-ice motion in a numerical model setting on one-day time scales.

Machine learning models for sea-ice have been applied to improve estimates of ice properties from satellite remote sensing (Lee et al., 2016; Dumitru et al., 2019), to predict and understand sea-ice concentration on different time scales (Kim et al., 2020; Li et al., 2021; Andersson et al., 2021), and to make predictions of sea-ice motion (Petrou and Tian, 2019; Zhai and Bitz, 2021). ML models have been successful at improving predictions of sea-ice properties in comparison to state-of-the-art dynamical models. For example, the deep learning model IceNet outperformed the SEAS5 dynamical model from the European Centre for Medium-Range Weather Forecasts (ECMWF) for lead times greater than one month when making seasonal forecasts of summer ice (Andersson et al., 2021). Additionally, a CNN designed to make one-day

predictions of ice motion showed higher correlations with satellite observations than CICE5, a leading physics-based model for sea ice (Zhai and Bitz, 2021). The high performance of this CNN provides evidence that a CNN would be an effective surrogate model to replace the sea ice dynamical component of a numerical model for short-time-scale predictions. We build upon the work of Zhai and Bitz (2021) by further analyzing the nuances in the performance of a CNN in predicting ice motion, and by building the case for its use over a conventional linear regression approach.

We apply three different models, including persistence (PS), linear regression (LR), and convolutional neural network (CNN) to make predictions of sea-ice motion. In comparison to the other two models, a CNN has the benefits of incorporating spatial information and non-linear relationships between the inputs into its predictions. We build a CNN that has a similar architecture to that of Zhai and Bitz (2021) (differences are noted in the supplementary information, Table B.1) and that is trained on the same input and output data. Our models show similar performances in making one-day predictions of sea-ice motion (Table B.1). We expand on previous work by putting an emphasis on understanding the spatial and temporal variability in performance of the different models and how it is related to various properties of the ice. We divide the Arctic into four geographic regions (Fig. 2.1) based on the differences in skill between the CNN and LR models, and we analyze model performance within each.

2.3 Background

Sea-ice motion, as described by the momentum equation (Equation 2.1), is determined from a balance of the momentum tendency ($\frac{D}{Dt}(m\vec{u})$) with drag from the atmosphere ($\vec{\tau}_a$) and ocean ($\vec{\tau}_w$), the Coriolis force ($mf\hat{k} \times \vec{u}$), the ocean surface tilt ($mg\nabla H$), and the internal ice stress ($\nabla \cdot \sigma$) (Olason and Notz, 2014; Feltham, 2008). The term on the left represents the total

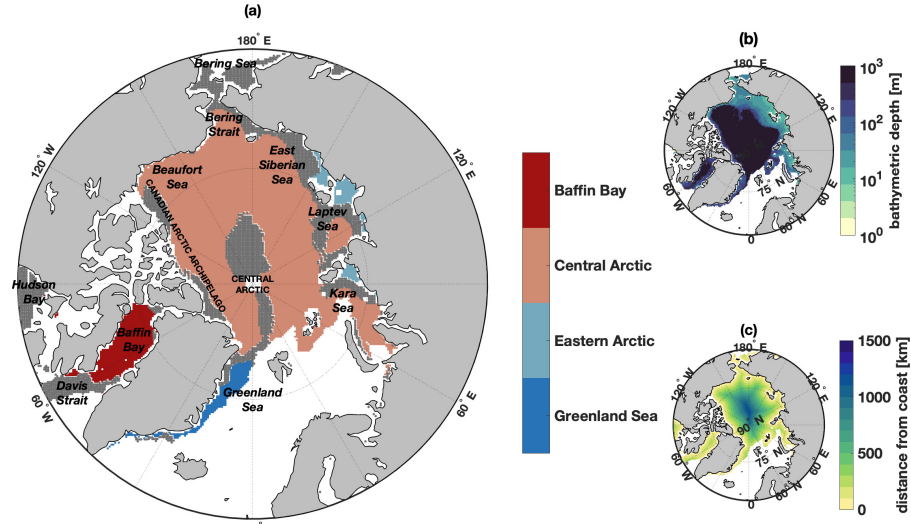


Figure 2.1. Maps showing (a) spatial divisions (Greenland Sea, Eastern Arctic, Central Arctic, and Baffin Bay), (b) bathymetric depth (m) (note logarithmic scaling), and (c) the distance from coast (km). Spatial divisions are based on overall performance of the CNN model and the difference between the performance of the CNN and LR models. The four divisions represent regions of: variable model performance and $corr_{LR} \gg corr_{CNN}$ (Greenland Sea, dark blue), low model performance and $corr_{LR} > corr_{CNN}$ (Eastern Arctic, light blue), high model performance and $corr_{LR} < corr_{CNN}$ (Central Arctic, light red), and variable model performance and $corr_{LR} \ll corr_{CNN}$ (Baffin Bay, dark red). Gray shading represents areas where the difference in correlation between the CNN and LR is not statistically significant or areas that are not included within this analysis. Data are not shown in regions where the ice concentration is zero or the satellite retrievals are absent for more than 20% of the year.

derivative of mass, m times velocity, \vec{u} :

$$\frac{D}{Dt}(m\vec{u}) = \vec{\tau}_a + \vec{\tau}_w - m f \hat{k} \times \vec{u} - mg \nabla H - \nabla \cdot \sigma. \quad (2.1)$$

Changes in external forcing (i.e. winds, currents, radiation, etc.) influence the geometric and mechanical properties of the ice (thickness distribution, mass, strength, drag coefficients, etc.), which ultimately impact ice motion and deformation (Untersteiner et al., 2007). The American-Canadian Arctic Ice Dynamics Joint Experiment (AIDJEX) of 1970-1978 was one of the first major studies aimed at developing a comprehensive model of sea-ice motion under the influences of the ocean and atmosphere (Maykut et al., 1972; Untersteiner et al., 2007). Using

data from the AIDJEX experiments, Thorndike and Colony (1982) introduced a relationship between sea-ice velocity and geostrophic wind that explained up to 70% of the variance in sea-ice velocity in the central Arctic. This relationship describes ice that is subject to high wind speeds on time scales of days to months. In this relationship, sea-ice velocity is related to geostrophic wind velocity through a speed reduction factor (the wind factor) and a turning angle, after removal of the long-term mean ice velocity field. In the absence of a steady ocean current, sea-ice moved about 8° to the right of the geostrophic wind at about 0.008 times the speed. This model is less successful for areas within 400 km of the coast, where stress gradients within the ice become more important due to the restriction of ice motion by geographical features (Thorndike and Colony, 1982).

The internal stress gradient also depends on factors including the magnitude of the wind speed, ice concentration, and ice thickness. Ice with high values for thickness and concentration may have large stress gradients, which can result in a smaller dependence on wind. Conversely, ice with smaller stress gradients (low thickness and concentration) is found to have higher dependencies on wind (Hibler, 1979). Decreases in correlation between wind and ice motion near the coast have often been attributed to ice stresses (Thorndike and Colony, 1982; Kimura and Wakatsuchi, 2000; Hibler, 1979).

A relationship between ice motion and geostrophic wind was also examined by Kimura and Wakatsuchi (2000) and by Maeda et al. (2020), using sea-ice motion derived from satellite products and geostrophic wind derived from the sea-level pressure data from ERA Interim Reanalysis data produced by European Centre for Medium-Range Weather Forecasts (ECMWF) on 2.5° and 0.75° grids, respectively. In these studies, geostrophic wind was generally found to explain 70% of the variance in sea-ice velocity, with 60–90% of the variance explained in the central Arctic, and up to 40% in coastal regions (Fig. 3 in Maeda et al. (2020)). In addition to spatial variability, seasonal variations in the speed reduction factor and turning angle have been reported (Thorndike and Colony, 1982; Kimura and Wakatsuchi, 2000; Kwok et al., 2013; Maeda et al., 2020).

2.4 Data

In our analysis, models are trained to make one-day predictions of sea-ice velocity given present-day wind velocity, previous-day sea-ice concentration, and previous-day sea-ice velocity from various satellite and reanalysis sources, during 1989–2021. Using present-day wind as a predictor of present-day sea-ice velocity incorporates information that gives the model intrinsic skill. This approach is appropriate for the objective to make predictions on one-day time scales. We opt not to detrend to avoid contaminating the data with spurious removals. However, we do find that the model performance does not have any significant changes when run on data with the seasonal cycle removed (not shown). Processed data and methods for obtaining and processing raw data are made available by (Hoffman et al., 2023).

The ice velocity and concentration data are available from 25 October 1978 to 31 December 2021. However, evaluation of the uncertainty metrics for the Polar Pathfinder ice motion product shows a change in the error fields starting in the summer of 1987 (Figure B.1) due to a difference in the sampling period when switching from using Scanning Multichannel Microwave Radiometer (SMMR, 48hr sampling period) to Special Sensor Microwave/Imagers (SSM/Is, 24hr sampling period) for brightness temperature (Tschudi et al., 2020). Additionally, ice concentration data from the Nimbus-7 passive microwave are only available every other day until 1987, and there is a gap in availability of the sea ice concentration data from 03 December 1987 to 12 January 1988. Thus, for consistency in the stability of the observation systems and the quantity of data used from each year, we use data from 1989–2021 to build our models. We use the satellite and reanalysis sources discussed below for consistency with Zhai and Bitz (2021). However, in comparison, we make a slight extension to the temporal subset of data over which the model is trained and tested.

2.4.1 Sea-Ice Velocity: Polar Pathfinder Version 4 Daily Sea Ice Motion vectors (PP)

The Polar Pathfinder product (PP; Tschudi et al., 2019) provides daily sea-ice motion vectors at a spatial resolution of 25 km in the Equal-Area Scalable Earth (EASE)-grid. The EASE-grid was defined by the NOAA/NASA Polar Pathfinder Program to support standardized spatial comparisons from gridded, satellite microwave data. In polar regions, the EASE-Grid takes the form of Lambert azimuthal equal-area projections that accurately represent area in all regions of the global sphere (Brodzik et al., 2012). This data set is informed by optimal interpolation of a combination of observations from passive microwave inputs, buoys, and NCEP/NCAR reanalysis winds. The PP dataset relies on wind because during the summer, passive microwave and buoy sources become unreliable for melting ice (Tschudi et al., 2020). For wind-derived ice motions, ice is assumed to move at $\sim 1\%$ of the wind speed and in the direction of the wind, based on the estimate from Thorndike and Colony (1982). An estimated uncertainty map is also provided, which we use for comparison when evaluating our models. We were unable to obtain a dataset that is independent from the PP product to validate the use of the PP for this case. We did find high correlation between the PP and the Ice-Tethered Profiler data (not shown), but these observations were used to create the PP product. Wang et al. (2022) found the PP to have low accuracy in speed, but high accuracy in angle in comparison to eleven other satellite products when evaluated against measurements from buoys from the International Arctic Buoy Program (IABP) and the Multidisciplinary drifting Observatory for the Study of Arctic Climate (MOSAiC).

2.4.2 Sea-Ice Concentration: Nimbus-7 SMMR and DMSP SSM/I-SSMIS Passive Microwave Data

The passive microwave sea-ice concentration product (Cavalieri et al., 1996) is generated from brightness temperature data derived from various sensors (SMMR, DMSP and SSM/I-SSMIS). This product provides daily measurements of sea-ice concentrations (fraction of ocean

area covered by sea ice in each grid cell) in a 25×25 km polar stereographic projection. Here we re-grid to the 25-km EASE-grid for consistency with other ML model inputs. An intercomparison study of 10 satellite passive microwave sea-ice concentration data sets by Kern et al. (2019) found that while the Nimbus-7 product used in this work showed the largest difference between other products, all 10 products compared reasonably well to ship-based observations. Additionally, the Nimbus-7 product used in this study showed less than a 7% deviation from all other products from November–June, and less than a 15% deviation from July–October when comparing the monthly mean values of sea-ice concentration among the 10 products from June 2002 to September 2011. The product used in this study was also found to have a negative bias in sea-ice concentration throughout the Arctic in comparison to the ensemble mean of the 10 products (Fig. 8 from Kern et al. (2019)). While this negative bias was particularly large in the peripheral seas, it was close to zero (i.e. $< 6\%$) in the region of study of this work.

2.4.3 Wind Velocity: Japanese 55-year Reanalysis derived for ocean-ice models (JRA55-do)

The Japanese Meteorological Agency 55-year atmospheric reanalysis based surface dataset for driving ocean-sea ice models (JRA55-do) is used to prescribe wind velocity (Tsujino et al., 2018). Based on the JRA55 (Kobayashi et al., 2015), the JRA55-do is derived for use in ocean simulations, with surface fields adjusted relative to satellite climatological winds (SSM/I and QuikSCAT) using a spatially varying wind factor for wind speed and EOF analysis for wind direction (Tsujino et al., 2018). The JRA55-do better matches satellite wind fields in coastal areas than do other reanalysis products (Taboda et al., 2019). The JRA55-do provides 3-hourly estimates of total wind velocity at 10 m with a horizontal resolution of ~ 55 km. Here we calculate daily average wind vectors and re-grid to the 25-km EASE-grid.

2.4.4 Bathymetric Depth: International Bathymetric Chart of the Arctic Ocean (IBCAO)

We use bathymetric depth from IBCAO (Jakobsson et al., 2020) for comparisons of model performance after training. We make use of the Version 4.2 product without elevation data for the Greenland Ice Sheet on a $400\text{ m} \times 400\text{ m}$ grid cell spacing, re-gridded to the 25-km EASE-grid.

2.5 Methods

2.5.1 Model Inputs

We employ a suite of machine learning and classical statistical models (PS, LR, and CNN) to predict present-day sea-ice velocity components ($u_{i,t}$ & $v_{i,t}$) using the input parameters:

- present-day zonal & meridional wind velocity ($u_{a,t}$ & $v_{a,t}$),
- previous-day zonal & meridional sea-ice velocity ($u_{i,t-1}$ & $v_{i,t-1}$), and
- previous-day sea-ice concentration (c_{t-1}).

Inputs are chosen based on results from Zhai and Bitz (2021), who showed that the above combination of parameters produced skillful output when used to predict sea-ice motion with a CNN. Sea-ice velocity might be expected to be dependent also on sea-ice thickness, in addition to our selected input fields (Hibler, 1979; Thorndike and Colony, 1982). However, feature exploration studies of CNN models applied to Community Earth System Model version 2 (CESM2) output by Zhai and Bitz (2021) found that the inclusion of sea-ice thickness as an input parameter does not greatly impact the overall skill and correlation of CNN predictions. Fortunately the thickness is not an important input, as satellite observations of sea-ice thickness prior to 2019 have a high uncertainty, are discontinuous in time and unavailable during the summer. Therefore this parameter is omitted from our analyses. We note that efforts are being made to extend the CryoSat-2 sea-ice thickness record back in time using machine learning

techniques (Landy et al., 2022). However, these data are available bi-weekly and thus do not meet the requirements of this study for daily data.

Inputs are taken from satellite and reanalysis sources listed in section 2.4. All variables are normalized to zero mean and one standard deviation before being input into the models, based on the global statistics of the entire record used here from 1989-2021. Data are broken up into train, validation and test data sets with an 88%–6%–6% split (e.g. train with years 1989–2017, validate with years 2018–2019 and test with years 2020–2021). The train, validate, and test years are shuffled ten times to produce data for ten different ensemble runs for each ML model. We refer to an “ensemble run” as a run that is trained on a different temporal subset of data. We calculate performance statistics (discussed in section 2.5.3) for each ensemble run and average over the ten runs for final results. A CNN requires inputs to be of consistent size, with consistent spatial and temporal coverage, and without non-numerical (e.g. not-a-number or ‘NaN’) values. Thus, while it may make sense to remove data in regions where sea-ice motion data are not available (i.e. sea-ice concentration is zero or there is land) before training, due to the practical constraints of applying a CNN sea-ice velocity components are set to zero during training. A time-variable mask is used to remove these sea ice free points during model evaluation. Additionally, while uncertainty metrics are available for the Polar Pathfinder sea-ice motion product, we do not mask out any points during training due to the constraints of CNN models listed above. We note that taking uncertainty into account during training of PS and LR models is possible, but to maintain consistency between models we leave that for future work.

2.5.2 Model Setup

We compare prediction outputs from three different models: PS, LR, and CNN.

Persistence and Linear Regression Models

PS predicts the present-day sea-ice velocity to be the same as the previous day at each grid point (Equation 2.2):

$$u_{i,t}^* = u_{i,t-1}^*. \quad (2.2)$$

This offers a baseline measure of the variability of the system and of the minimum skill that any alternative models should attain. Here the vector u_i^* is a complex number, where the real and imaginary parts are the zonal and meridional components of the sea-ice velocity vector.

LR regresses each of the five input parameters (section 2.5.1) onto the sea-ice velocity components (Equation 2.3):

$$u_{i,t}^* = Au_{a,t}^* + Bu_{i,t-1}^* + Cc_{i,t-1}^* + D \quad (2.3)$$

Given inputs and outputs, LR solves for parameters A to D . In equation (2.3), A to D are complex constants, and the vectors u_i^* , u_a^* and c_i^* are complex numbers, where the imaginary part of c_i^* is set to zero. LR is carried out in two different manners: one is performed globally (LR-g), and uses each time snapshot as an independent sample for fitting, providing one equation for the entire modeled region in the Arctic; the other is performed grid-wise (LR), leading to a different regression equation for each grid point. For both LR configurations we employ ridge regression with a ridge parameter of $\lambda = 10^{-2}$ to limit the magnitude of the regression coefficients and prevent them from being unrealistically large (Marquardt and Snee, 1975). The value of the ridge parameter is chosen based on the iterative approach in Marquardt and Snee (1975) where we make step changes from small to large values of λ and pick the value of λ for which the LR coefficients stabilize (i.e. are not infinitely large). We also note that data are not removed from the training set when $c_i = 0$, which may dampen the wind dependence in LR because the model is trained that $u_i = 0$ when $u_a \neq 0$ in these locations. As discussed in section 2.5.1, these data are not masked during training because the CNN requires numerical values (i.e. not ‘NaN’).

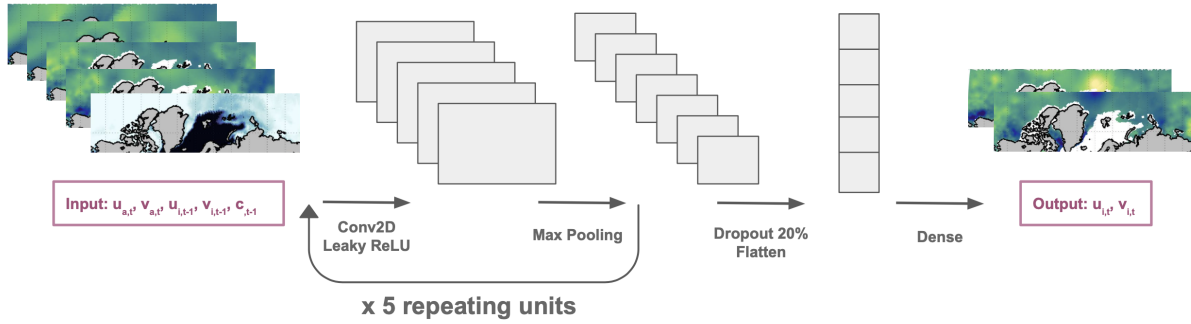


Figure 2.2. Schematic of the convolutional neural network (CNN) used in this study for predicting present-day sea-ice velocity components, $u_{i,t}$ & $v_{i,t}$ (outputs), from present-day wind velocity, $u_{a,t}$ & $v_{a,t}$, previous-day sea-ice velocity, $u_{i,t-1}$ & $v_{i,t-1}$, and previous-day sea-ice concentration, $c_{i,t-1}$ (inputs). This CNN has five repeating units of a 2D convolution with a ReLU activation and max pooling, followed by a 20% dropout layer, flattening and a dense layer.

Convolutional Neural Network (CNN) Architecture

A CNN architecture based on Zhai and Bitz (2021) is used in this study. A CNN is a type of ML model typically applied to visual images, where a computer is fed numerous (hundreds to millions) different images and learns from their patterns in order to make a prediction (O’Shea and Nash, 2015). We use data sets that are image-like in that they have a specified value at various grid-points on a map (for images this would be colors at various pixel locations). Incorporation of spatial information when making predictions is one of the benefits of CNN over LR or PS models, in addition to the ability of a CNN to capture non-linearities in the relationships between the input predictors and the outputs. Our CNN (Fig. 2.2) is set up with five repetitions of the block unit: 2D convolution, ReLU (Rectified Linear Unit), and 2D max-pooling. This is followed by a 20% drop-out layer, a flattening to a one-dimensional vector, and finally a regression onto a 1D vector (dense layer) representative of the output predictions. This output is then concatenated into two maps of present-day zonal and meridional sea-ice velocity.

We implement the CNN in python using the Tensorflow/Keras library (Abadi et al., 2015). Convolutional and ReLU layers are carried out with (1,1) strides and (3,3) filter sizes, whereas the max pooling strides and filter sizes are (2,2). For each of the respective repeating block units, there are 7, 14, 28, 56, and 112 filters. The training runs for 50 epochs with a batch size of 365

days. Optimization is carried out with an Adam optimizer and a normalized root mean squared error as the loss function (second term in Equation 2.5 discussed below). Similarly to the LR, we employ ridge regression with a ridge parameter of $\lambda = 10^{-2}$. Further descriptions of the architectural components of a CNN (i.e. layers, strides, filters, ReLU, max pooling, etc.) can be found in O’Shea and Nash (2015). Filter sizes are chosen based on the conventional VGGNet architecture (Szegedy et al., 2015). We do not carry out hyperparameter tuning for this study in order to maintain consistency with the architecture of Zhai and Bitz (2021), with the only differences being in the sizes and number of the filters due to differences in the sizes of the starting input maps.

2.5.3 Model Evaluation

As in Zhai and Bitz (2021), the model performances are evaluated and compared based on the correlation (Corr) and skill, given by:

$$corr_{x,y} = \frac{\sum_i^n (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i^n (x_i - \bar{x})^2} \sqrt{\sum_i^n (y_i - \bar{y})^2}}, \quad (2.4)$$

$$skill_{x,y} = 1 - \frac{\sqrt{(x_i - y_i)^2}}{\sqrt{(x_i - \bar{x})^2}}, \quad (2.5)$$

where x represents observations, and y represents predicted values of a sample size n . The correlation (Equation 2.4) is defined as the covariance between prediction and observation scaled by their standard deviations. The skill (Equation 2.5) is a representation of the fraction of the observed standard deviation explained by the model predictions, where the second term is the root mean squared error normalized by the standard deviation of the observations (Thomson and Emery, 2014). The correlation ranges from -1 to 1, with 1 indicating a perfect positive relationship, -1 indicating a perfect negative relationship, and zero representing orthogonality. The skill can range from negative infinity to 1, with 1 representing a perfect match between model

predictions and observations. The correlation is a measure of how well the phase variability in the data is explained by the model, whereas the skill is a measure of the absolute error in the model predictions.

These metrics are calculated using the test data set (varying years, as discussed in section 2.5.1) of which the models have no prior knowledge. Two different masks are made and both applied to the data during model evaluation: one is time-variable and evaluates model outputs only at times and in locations where sea-ice concentration is greater than zero; the other is constant with time and masks out all areas where sea-ice concentration is zero more than 20% of the time from 1992–2017. Metrics are calculated overall (section 2.6.1), at each grid point to provide spatial evaluation (section 2.6.1), over each month for temporal evaluation (section 2.6.1), and for different percentile ranges of various sea-ice properties (wind speed, ice speed, and ice concentration) to understand the role these play on the model performance (section 2.6.1). For temporal evaluations we calculate the monthly mean for each of the ten ensemble runs. Overall monthly means are then represented by the mean of the ten ensemble runs, and monthly errors are calculated as the standard error of the mean of the ten ensemble runs (as discussed in section 2.5.1). Temporal evaluations are carried out for different regions within the Arctic. The divisions (Fig. 2.1a) are made based on spatial distributions of model performance metrics ($corr_{CNN}$ and $corr_{CNN} - corr_{LR}$ in Fig. 2.3c & f), representing regions of: (i) variable model performance and LR greatly outperforming CNN (i.e. variable $corr_{CNN}$ and $corr_{LR} \gg corr_{CNN}$; Greenland Sea, dark blue), (ii) low model performance and LR slightly outperforming CNN (i.e. low $corr_{CNN}$ and $corr_{LR} > corr_{CNN}$; Eastern Arctic, light blue), (iii) high model performance and CNN slightly outperforming LR (i.e. high $corr_{CNN}$ and $corr_{LR} < corr_{CNN}$; Central Arctic, light red), and (iv) variable model performance and CNN greatly outperforming LR (i.e. variable $corr_{CNN}$ and $corr_{LR} \ll corr_{CNN}$; Baffin Bay, dark red).

2.5.4 Model Comparison

We also investigate the correlation and skill differences between the LR and CNN models, which requires an understanding of where the differences are significant. Significance tests on the differences are approximated with a cross-validated t test (Dietterich, 1998; Tang et al., 2000). The cross-validated t test proceeds as follows: (i) for each of the ten ensemble runs, the correlation and skill for the LR and CNN are calculated for each grid point or percentile range for a given variable and transformed by Fisher’s z transform (Equation 14.5.6 in Press et al., 1986) to remove skewness in the distribution; (ii) the difference between the transformed correlation and skill for the two models is calculated and averaged over the ten ensemble runs; (iii) a two-tailed t test is performed to detect whether the mean difference between the two models is significantly different from zero at the 95% confidence level. The cross-validated t test uses the degrees of freedom to calculate significance. For spatial comparisons (section 2.6.1) we estimate degrees of freedom using the temporal decorrelation scale to estimate the number of independent time series of sea-ice motion in the Arctic. This temporal decorrelation scale is taken as the e -folding scale of a Gaussian fit to the autocorrelation of the sea-ice speed calculated at different time-lags (Equations 10 and 11 in Sumata et al., 2018).

2.5.5 Analysis of Inputs

We analyze the spatial and temporal variability of different parameters related to ice motion (wind speed, u_a ; ice speed, u_i ; and ice concentration, c_i) to assess how the model performance compares to the model inputs. Spatial analyses look at maps of the average and standard deviation of each parameter over time from 1989–2021. This type of analysis is useful for comparing these properties to maps of the model performance metrics in order to understand different regimes within the Arctic. We also look at the seasonality of each of these properties. Similarly to the input analysis in section 2.5.3, monthly errors are calculated as the standard error of the mean of the ten ensemble runs, and temporal evaluations are carried out for different

divisions that are chosen based on the model performances.

2.6 Results

2.6.1 Model Performance

Overall

We evaluate the overall performance of the different models by calculating the correlation and skill over all gridpoints and times (Table 2.1). The CNN has the highest correlation and skill, followed closely by the grid-wise linear regression (LR). The grid-wise LR largely outperforms the global LR (LR-g) that covers the entire Arctic, which is not much better than the simple PS model. These results confirm the advantage of using a model that captures non-linearity (CNN) and the heterogeneity of Arctic sea ice motion statistics (both CNN and LR). The better performance of the CNN, LR, and LR-g models in comparison to the PS confirms that sea-ice motion depends on wind and sea-ice concentration on daily time scales. Table 2.1 shows the pattern that an increase in model complexity leads to an increase in performance. Additionally, because correlation is a measure of how well the model is able to capture the phasing, while skill measures the model’s ability to capture phasing and magnitude, the high correlation but lower skill suggests the models do well capturing the phasing but incur error in capturing the magnitude.

Table 2.1. Overall correlation and skill between observations and predictions of sea-ice velocity for four different models.

Model	Correlation	Skill
Persistence (PS)	0.69 ± 0.02	0.21 ± 0.02
Linear Regression, global (LR-g)	0.72 ± 0.01	0.30 ± 0.01
Linear Regression, gridwise (LR)	0.78 ± 0.02	0.37 ± 0.02
Convolutional Neural Network (CNN)	0.81 ± 0.01	0.42 ± 0.02

Spatial

Spatial variations in the correlation (Fig. 2.3) and skill (not shown) are similar for the PS, LR, and CNN models. Models perform well for predictions in the central Arctic, with decreasing performance in coastal locations. Low values of correlation (Fig. 2.3a–c) are visible in the Bering Strait, Bering Sea, Hudson Bay, East Siberian Sea, Laptev Sea, Kara Sea, and off the coast of Greenland. Particularly poor model performances are found near the islands in the Eastern Arctic. The best model performance is seen north of Fram Strait and in the Beaufort Sea.

Typically, $corr_{CNN} > corr_{LR} > corr_{PS}$, similar to the results from section 2.6.1. The spatial differences in correlation between the models are shown in Fig. 2.3d–f. Regions in red indicate areas where the first model in the difference metric outperforms the second (i.e. $corr_{CNN} > corr_{PS}$ in Fig. 2.3d, $corr_{LR} > corr_{PS}$ in Fig. 2.3e, and $corr_{CNN} > corr_{LR}$ in Fig. 2.3f), whereas blue regions indicate the opposite (i.e. $corr_{CNN} < corr_{LR}$ in Fig. 2.3f). Gray regions show where the difference in correlation between the two models is not statistically significant. The CNN and LR models outperform the PS over the entire Arctic (Fig. 2.3d & e), with the exception of the western side of Baffin Bay where the PS outperforms the LR (blue). Overwhelmingly, the CNN outperforms the LR (red in Fig. 2.3f). Interestingly, the LR has a higher correlation (blue) in coastal regions where both models have decreased performance (i.e. near the islands in the Eastern Arctic and off the coast of Greenland).

The spatial patterns in model performance compared to the distance from the coast are confirmed in Fig. 2.4. Correlations for the CNN and LR models tend to be lower in coastal regions (Fig. 2.4a–b). This is also true for skill (not shown). For both models, locations that are greater than 400 km from the coast consistently have correlation greater than 0.7 (and skill greater than 0.3, not shown). The finding that the CNN outperforms the LR model for most cases is confirmed in Fig. 2.4c, where most of the data lie in the positive region (i.e. above the black line). Conversely, locations where the LR outperforms the CNN only occur within 400 km of the coast.

We also show that models have decreased performance in shallower regions (Fig. 2.4d–e). Overall, model performance increases with increasing seafloor depth. The relationship is logarithmic: performance increases rapidly with increasing depth for depths shallower than 1000 m, while the trend levels out for depths greater than 1000 m. Models exhibit correlations less than 0.7 and 0.5 (CNN and LR, respectively) only for locations with depths less than 1000 m. The CNN outperforms the LR for most cases (Fig. 2.4f). Most regions where the LR outperforms the CNN (below the black line) occur at depths shallower than 500 m, although there are some instances of higher correlation of the LR for greater depths.

We also analyze the spatial variability of the various properties related to sea-ice motion (wind speed, u_a ; sea-ice speed, u_i ; and sea-ice concentration, c_i). The mean and standard deviation of the properties listed above are mapped in Fig. 2.3g–l. Patterns in mean ice speeds tend to coincide with the spatial patterns in wind speed (Fig. 2.3g & h), consistent with the known dependence of ice motion on wind speed (Thorndike and Colony, 1982). Both ice and wind speed are relatively low in the coastal and island regions of the East Siberian Sea, the Canadian Arctic Archipelago, and off the northern and western coasts of Greenland. The highest mean wind speeds occur in the Davis Strait, off the eastern coast of Greenland, and in the Bering Strait; high mean ice speeds also occur in these regions, in addition to the Beaufort Sea. The region of low mean ice speeds to the north of the Canadian Arctic Archipelago coincides with high mean ice concentrations (Fig. 2.3h & i). Conversely, the region of low mean ice speeds in the East Siberian Sea coincides with lower mean ice concentrations.

Regions that show high variability (large standard deviations) in ice speed coincide with high mean ice speeds (i.e. in the Beaufort Sea, Baffin Bay, Davis Strait, and Greenland Sea), while regions with low variability coincide with lower mean ice speeds (to the north of the Canadian Arctic Archipelago and in the East Siberian sea) (Fig. 2.3h & k). Variability in wind speed is found to be relatively consistent throughout the Arctic, with the exception of high variability off the eastern coast of Greenland (Fig. 2.3j). Regions with large variability in ice concentration typically correspond to regions with lower mean ice concentrations (i.e. in the

East Siberian Sea, Baffin Bay, the Kara Sea, and the Bering Strait). These are the regions where the largest amount of seasonal ice melt typically occurs (not shown), which contributes to the large variability and lower mean ice concentrations.

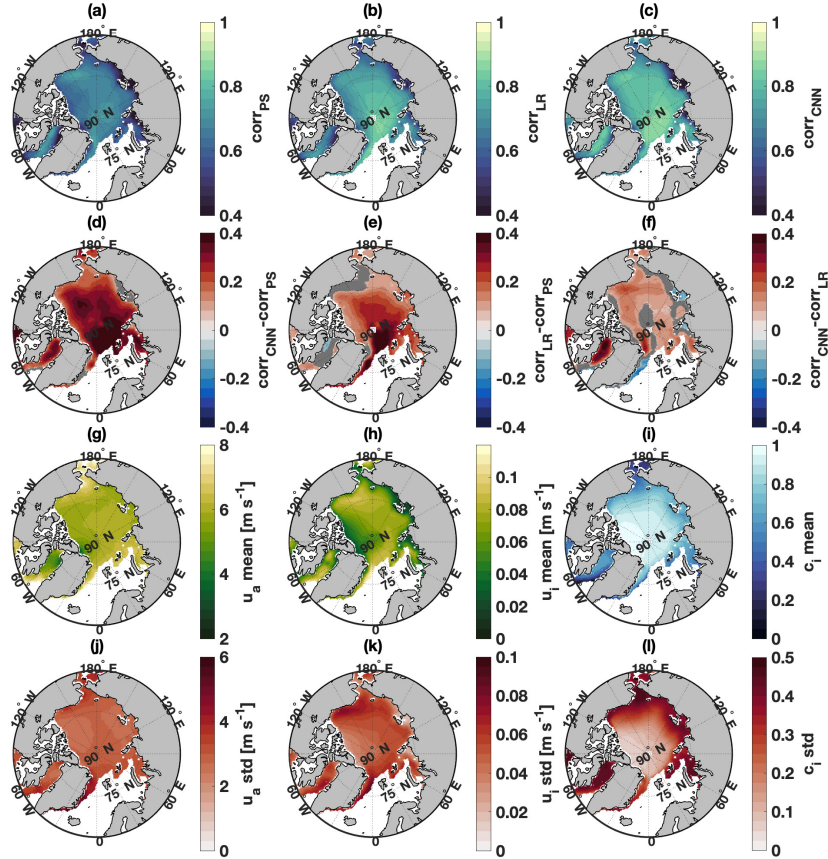


Figure 2.3. (a–c) Mapped correlation for predictions of sea-ice velocity made by the (a) PS, (b) LR, and (c) CNN models. (d–f) The difference in correlation between models: (d) $corr_{CNN} - corr_{PS}$, (e) $corr_{LR} - corr_{PS}$, and (f) $corr_{CNN} - corr_{LR}$. The gray regions in d–f represent locations where the difference in correlation between the two models is not statistically significant. (g–i) Mean and (j–l) standard deviation in time of various properties related to sea-ice motion from satellite and reanalysis products (wind speed, u_a (g & j); sea-ice speed, u_i (h & k); and sea-ice concentration, c_i (i & l)).

Temporal

For the region containing the entire Arctic, the CNN typically has the highest correlation, followed by the LR and then the PS model (Fig. 2.5a). During June–September the difference in correlation between the CNN and LR models is not statistically significant. Temporal structure

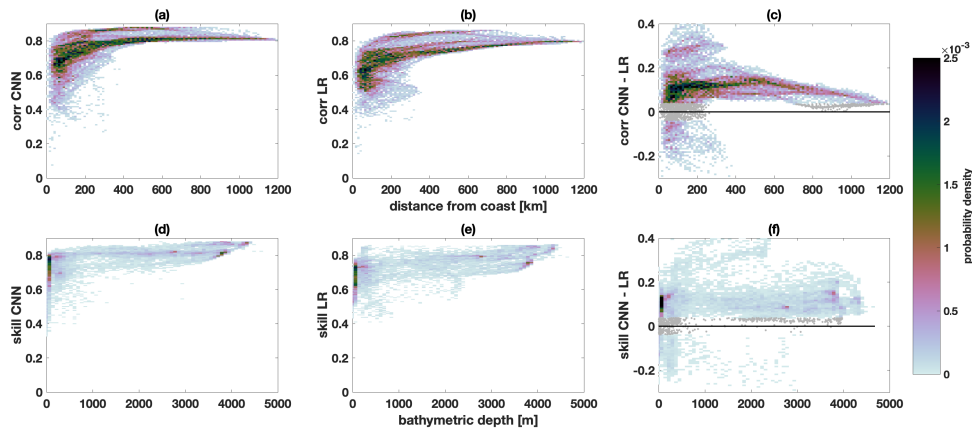


Figure 2.4. PDFs for model performances compared to their distance from the coast (a–c) and bathymetric depth (d–f), with (a & d) for the CNN and (b & e) for the LR. (c & f) The difference ($corr_{CNN} - corr_{LR}$) between the correlation of the two models. Gray shading (c & f) represents correlation differences between the two models that are not statistically significant. Results for skill (not shown) are similar to correlation.

is visible in the correlation for all of the models. The LR model performance (Fig. 2.5a) has a larger range of seasonal variability than the other two models. Maximum correlation and skill for the PS and CNN models occurs during October–December, while the LR has a correlation maximum in June–August. All three models experience a minimum performance in April.

The temporal evaluations are divided into regions (Fig. 2.1a) based on the spatial variability of their performance, as discussed in section 2.5.3. The impacts of this spatial division on model performance are shown in Fig. 2.5b–d, while Fig. 2.5e–g represent differences in the correlation between the different models. Here the black lines represent metrics calculated with all of the data included, and the different shades of red and blue represent the respective spatial regions from Fig. 2.1a. Diamonds in Fig. 2.5e–g indicate months where the difference between the two models is statistically significant. The correlation for the region within the Central Arctic division (light red) does not deviate much from that of the entire Arctic (black) because the Central Arctic region is large and covers most of the region containing the entire Arctic. However, there are significant changes in monthly values of correlation for all other divisions (Greenland Sea, Eastern Arctic, and Baffin Bay divisions). For all three models, the Eastern

Arctic (light blue) division exhibits a similar seasonal cycle to the entire Arctic (i.e. minimum correlation in March–April), but has a consistently lower monthly correlation in comparison to the other divisions for all models, except during the months of July–October.

The Greenland Sea (dark blue) and Baffin Bay (dark red) divisions exhibit a relatively high correlation from October–May that decreases toward a minimum in August or September (Fig. 2.5b–d). The Greenland Sea division (dark blue) has a higher correlation than the other divisions from October–April for all three models. The Greenland Sea division shows a lower correlation than the region containing the entire Arctic from the months of June–September, reaching a minimum in August for all three models that is significantly lower than correlations for the entire Arctic (i.e. the CNN has a minimum of 0.54 for the Greenland Sea division in comparison to 0.80 for the overall Arctic). The Baffin Bay division (dark red) exhibits the largest deviations in correlation from the overall Arctic for all models, showing up as a large decrease during the months of May–November. The Baffin Bay division has higher correlations in December–April, and the lowest August–September minimum out of all of the divisions for all models (i.e. the August correlation of the CNN within the Baffin Bay division is 0.28 in comparison to 0.80 for the entire Arctic). The performance minima that occur in August–September for the Greenland Sea and Baffin Bay regions are much lower than the April minima for the region containing the entire Arctic. This pattern of decreased model performance during months of minimum sea-ice extent (Greenland Sea and Baffin Bay divisions) suggests a link between model performance and sea-ice concentration, which will be further evaluated in section 2.6.3.

The differences in correlation between the models for the different divisions are shown in Fig. 2.5e–g. The LR and CNN typically outperform the PS for all divisions (i.e. diamonds indicating statistically significant difference in model performance are above zero in Fig. 2.5e & f). The LR outperforms the CNN in all months for the Eastern Arctic division. However, statistically significant differences from zero are only present December–May. The CNN outperforms the LR during the months of September–May for the Central Arctic, and September–June for the

Baffin Bay division. However, the difference between the correlation of the CNN and LR is not statistically significant during the months of June–October for the Central Arctic or July and September–November for the Baffin Bay division. These differences in model correlation will be further analyzed in section 2.6.3.

We also compare the temporal variability in performance to that of the various properties related to sea-ice motion (wind speed, u_a ; sea-ice speed, u_i ; and sea-ice concentration, c_i). The ensemble mean monthly averages of various properties related to sea-ice motion are shown in Fig. 2.5h–j. Analysis is further broken down into the four divisions within the Arctic, which are chosen based on values of the model correlations (Fig. 2.1).

For all regions the seasonal cycles for ice speed and wind speed (Fig. 2.5h & i) generally line up, with minima typically occurring during the summer months and maxima in the winter. The seasonal pattern of minimum wind speeds occurring from June–July, and maximum speeds anywhere from October–February is consistent throughout all regions, except for the Eastern Arctic division where minima are found in December–March, and maxima occur in September–October. The Greenland Sea division has greater seasonal variability in wind and ice speeds than the other divisions, with comparatively high maximum speeds in November–April. Seasonal patterns in ice speed show minima in June–July for the Central Arctic, June–August for the Greenland Sea, and May–October for the Baffin Bay division. The Eastern Arctic division shows the opposite seasonal trend, instead exhibiting minimum ice speed from December–May. Sea-ice concentration also follows a seasonal cycle within each division, typically reaching a maximum in March and a minimum in September (Fig. 2.5j). The Baffin Bay division exhibits the lowest and longest duration minimum ice concentration (i.e. $c_i < 0.5$ from July–October). From December–May the Greenland Sea division has a lower ice concentration than the other divisions, which are all similar during this time.

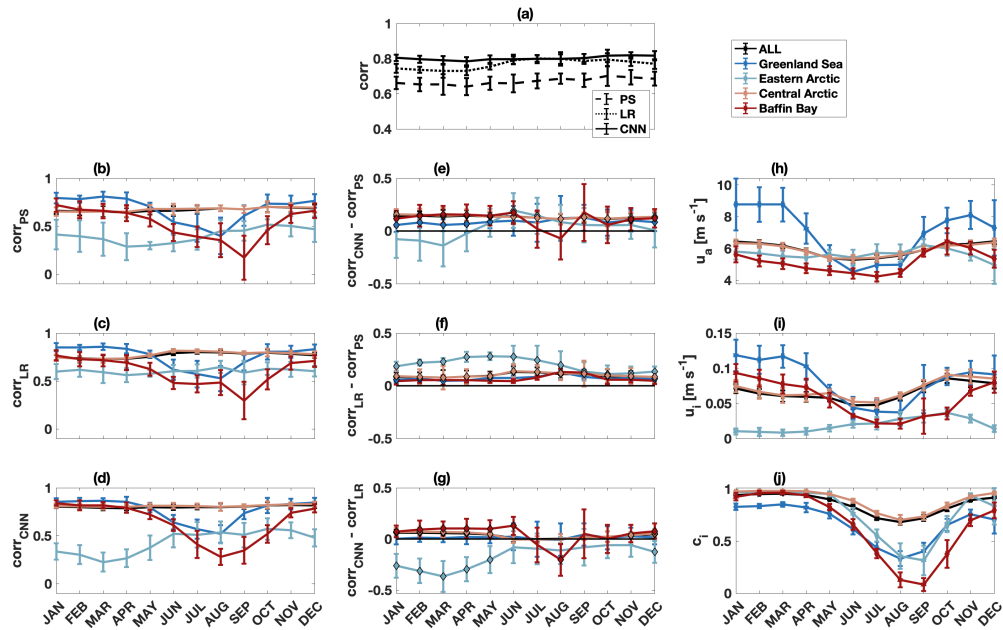


Figure 2.5. (a–d) Ensemble mean monthly correlation for the prediction of sea-ice velocity by three different models: (a) all models, (b) persistence, (c) linear regression, and (d) CNN. (e–g) The difference between the correlation of the (e) CNN and PS, (f) LR and PS, and (g) CNN and LR models. (h–j) Ensemble mean monthly values of various properties related to sea-ice motion (wind speed, u_a (h); sea-ice speed, u_i (i); and sea-ice concentration, c_i (j)). Metrics are calculated for five different regions: containing the entire area of the Arctic (black), and within the spatial divisions indicated in Fig. 2.1a (shades of red and blue). Error bars represent ensemble mean standard deviations. Diamonds in Fig. 2.5e–g indicate months where the difference between the two models represented is statistically significant.

Model Performance for Percentiles of Inputs

The model performance is compared to properties related to sea-ice motion (wind speed, u_a ; sea-ice speed, u_i ; and sea-ice concentration, c_i) to probe the variability in model correlation in space and time. Figure 2.6 shows the correlation metrics calculated from subsets of test data for all models (PS in dark blue, LR in teal, and CNN in green). Subdivisions are based on percentile ranges (5% intervals) of the various properties. The performance metrics (correlation (a–c) and the difference in correlation between the various models (d–f)) are plotted against the average of each percentile range (i.e. 0–5%, 5–10%, etc.) for each property. Skill metrics (not shown) have

similar patterns to the correlation. We find that the correlation increases with increasing wind speed, sea-ice speed, and sea-ice concentration for all models (Fig. 2.6a–c). These relationships have statistically significant r^2 values when fit to a second-order polynomial with a least squares regression.

The CNN and LR consistently outperform the PS model, as these two difference metrics ($corr_{CNN-corr_{PS}}$ & $corr_{LR-corr_{PS}}$) are positive for all u_a , u_i , and c_i (blue and teal lines in Fig. 2.6 d–f). The CNN has a higher correlation than the LR (green lines in Fig. 2.6 d–f), except for the case where $c_i < 0.5$ (Fig. 2.6f). The metrics for the difference between the CNN and the other two models (i.e. $corr_{CNN-corr_{PS}}$ & $corr_{CNN-corr_{LR}}$) have statistically significant relationships with wind speed, ice speed, and ice concentration: the difference between the two models decreases for increases in wind and ice speed (Fig. 2.6d & e), and increases with increases in ice concentration (Fig. 2.6f). The difference metric $corr_{LR-corr_{PS}}$ shows a similar relationship to u_i , but not u_a or c_i . Additionally, the difference between the CNN and the LR is less dependent on u_i than the other two difference metrics (i.e. the slope of the green line is less than the slopes of the teal and blue lines in Fig. 2.6e). This can be attributed to the correlation of the PS model being much lower than that of the CNN or LR when ice speeds are close to 0 m s^{-1} . The results in Fig. 2.6d–f are robust whether we use all data or remove non-significant points.

2.6.2 Linear Regression Parameters: Relationship Between Sea-Ice Motion and Input Parameters

Analysis of the linear regression parameters provides insight on the locations where each of the inputs is important for predicting sea-ice motion. The parameters from the full LR (A – C in Equation 2.3) described in section 2.5.2 are mapped in Fig. 2.7. Here Fig. 2.7a–c represents the magnitude of the regression coefficients for normalized wind speed, sea-ice speed, and sea-ice concentration on the sea-ice velocity (i.e. $\sqrt{\Re^2 + \Im^2}$ of A to C , where \Re and \Im represent the real and imaginary components of these coefficients). These values range between 0 and 1 in the figure because they are normalized to the maximum overall coefficient. Larger values indicate

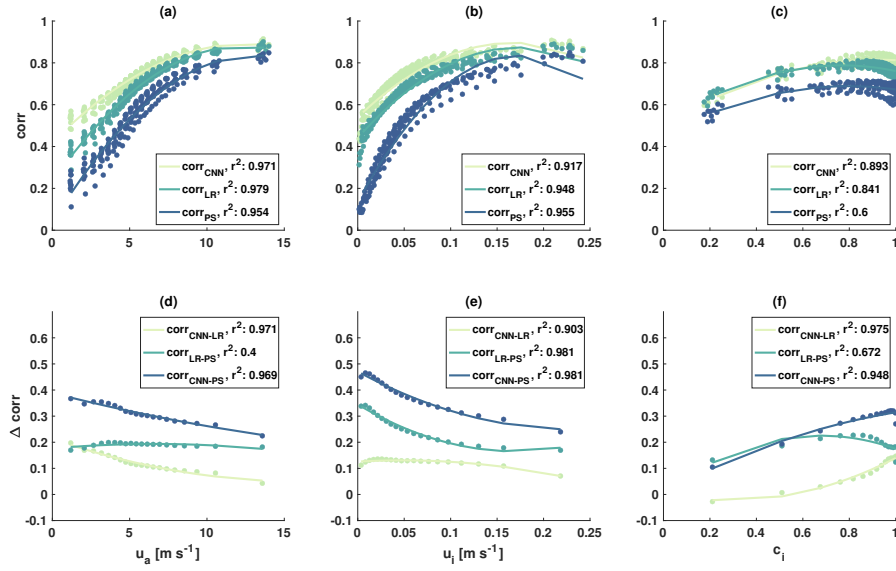


Figure 2.6. Correlation of the CNN (a–c) and the difference between CNN and LR correlation (d–f) as a function of various properties related to sea-ice motion (wind speed, u_a (a & d); sea-ice speed, u_i (b & e); and sea-ice concentration, c_i (c & f)). The correlation is calculated with subsets of test data based on percentiles (5 percent intervals) of the various parameters. The x-axis represents the mean value of the data in each 5% interval of each parameter. Correlation differences (d–f) that are not statistically significant are not shown.

that sea-ice velocity has a larger linear dependence on a particular parameter.

Results show that wind speed has the largest importance in predicting sea-ice velocity within the Central Arctic (Fig. 2.7a). Near the coast, the LR coefficient for previous-day sea-ice velocity is elevated (Fig. 2.7c) complementary to the high values in the interior for wind speed (Fig. 2.7a). Fig. 2.7d–e represents the rotation angles of the wind and sea-ice velocity to the predicted next-day sea-ice velocity. The wind angle has an average of $24.9^\circ \pm 11.3^\circ$ throughout the Arctic, which is fully consistent with Nansen’s observations aboard the Fram of angles between 20 and 40° (Ekman, 1905), and falls within one standard deviation of previous research (Thorndike and Colony, 1982; Serreze et al., 1989; Maeda et al., 2020) who found wind angles of -5 to 18° , 0 to 19° , and -10 to 30° (depending on season; winter to summer), respectively. The spatially averaged angle between present and previous-day sea-ice speed is $-8.3^\circ \pm 6.4^\circ$, with spatial variations as seen in Fig. 2.7e. When looking at the data, the expected spatial mean of the

angle difference between previous and present-day sea-ice velocity is 0.2° (not shown), which is within two standard deviations of the angle found from the LR parameters.

Wind velocity is found to have the maximum LR coefficient for predicting sea-ice velocity throughout the Central Arctic (dark blue in Fig. 2.7f). Locations near the coast are dominated by the sea-ice speed (pink regions). This is consistent with results from previous studies (Thorndike and Colony, 1982; Kimura and Wakatsuchi, 2000; Maeda et al., 2020) that conclude that the dependence of sea-ice velocity on wind velocity is not as strong in coastal locations where ice stresses become more important. Additionally, the low coefficient for wind velocity found in the Fram Strait off the east coast of Greenland, where the transpolar drift acts as a strong and persistent export pathway for Arctic sea-ice (Weiss, 2013), has previously been attributed to strong surface ocean currents (Kimura and Wakatsuchi, 2000).

The LR coefficient for wind speed is related to the spatial patterns in the mean c_i (Figs. 2.7a & 2.3i). We find low values for the LR parameter for wind speed in the Canadian Arctic Archipelago, a region where c_i is high and has little temporal variability (Fig. 2.3i & 1), which is consistent with results from Kimura and Wakatsuchi (2000); Maeda et al. (2020). However, regions of low mean c_i often have smaller values for the LR wind coefficient (i.e. coastal regions in the eastern Arctic, Baffin Bay, and the Bering Strait). This contradicts results from Kimura and Wakatsuchi (2000); Maeda et al. (2020), where areas with high ice concentration exhibit a relatively small wind factor as a result of internal stresses becoming more important in regions where ice is thick and concentrated. However, we note that in contrast to Kimura and Wakatsuchi (2000), our model also includes u_i as a predictor, which increases in importance near the coast. Additionally, our analysis has one LR coefficient at each spatial location throughout all time from 1992–2017, which provides a description of the relationship between the wind factor and the average c_i at each location. In contrast, Maeda et al. (2020) have a different LR equation for each month, providing a better picture of the relationship between the wind factor and the instantaneous c_i , which is more likely to display impacts of ice stresses.

Values of the LR coefficients are related to the performance of the LR model and to the

difference between the CNN and LR model performance. Figure 2.8 shows the relationship between the LR coefficients and the model correlation (Fig. 2.7a–c), and the difference between the correlation of the CNN and the LR (Fig. 2.7d–e), as calculated at each grid point. Larger LR coefficients for wind speed are associated with larger correlation of the LR model (Fig. 2.8a) in addition to an improved performance of the CNN over the LR (Fig. 2.8d). Conversely, a larger LR coefficient for sea-ice speed is associated with lower correlation (Fig. 2.8b) and does not show a statistically significant relationship with the difference metric, $corr_{CNN} - corr_{LR}$ (Fig. 2.8e). A larger LR parameter for ice concentration is linked to higher model correlation (Fig. 2.8c) and tends toward the LR outperforming the CNN (Fig. 2.8f). The skill (not shown) exhibits the same patterns as the correlation.

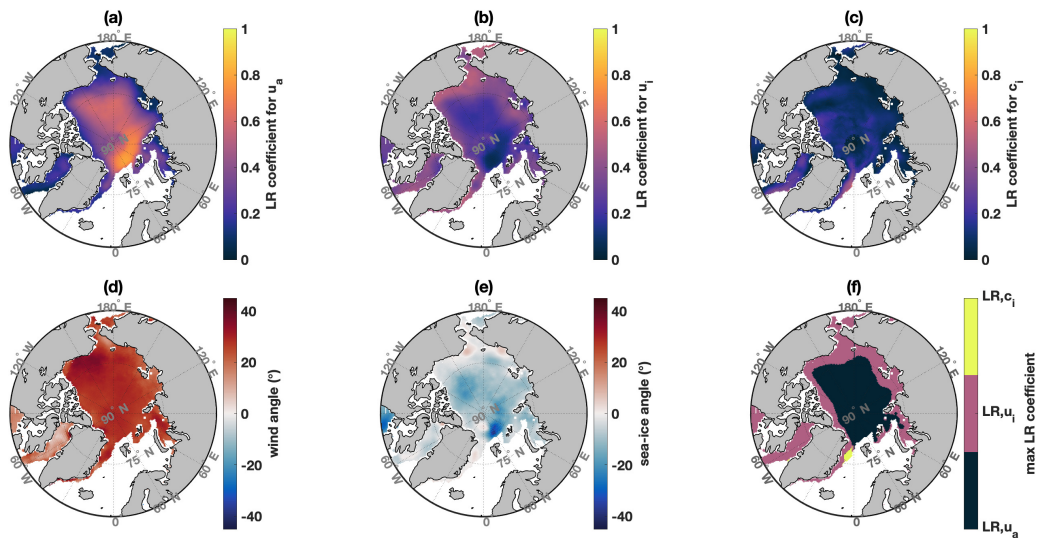


Figure 2.7. (a–c) Magnitude of the normalized linear regression coefficient for the relationship between sea-ice velocity components and input parameters (a, wind speed, A ; b, sea-ice speed, B ; c, sea-ice concentration, C) normalized to the maximum of a–c. (d–e) Mean angle of (d) wind speed and (e) sea-ice speed to the predicted next-day sea-ice speed. (f) Maximum linear regression parameter (a–c) for predicting sea-ice velocity at each location. Wind and ice speed parameters are derived from calculating the magnitude of the parameters for the velocity components.




Figure 2.8. PDFs for LR correlation (a–c) and the difference between the correlation of the CNN & LR correlation (d–f) compared to the LR coefficient magnitudes for (a & d) wind speed, LRu_a ; (b & e) sea-ice speed, LRu_i ; and (c & f) sea-ice concentration, LRc_i .

2.6.3 Attribution assessment of model predictive skill

We address our aims to understand (i) reductions in forecast skill and (ii) discrepancies in the performance of the different models by comparing the variability of these performance metrics (i.e. $corr_{CNN}$ and $corr_{CNN} - corr_{LR}$) to variables related to ice motion (i.e. distance from coast, d_c ; bathymetric depth, d ; wind speed, u_a ; ice speed, u_i ; ice concentration, c_i ; and the LR coefficients for wind speed, A , ice speed, B , and ice concentration, C). We focus on the difference between the CNN and the LR, because the CNN and LR both outperform the PS for almost all spatial locations.

In section 2.6 we find high model performance is linked to large distances from the coast, depths (Fig. 2.4 in section 2.6.1), wind speed, ice speed, ice concentration, (Fig. 2.6 in section 2.6.1), and values of the LR coefficients for wind speed & ice concentration (Fig. 2.8 in section 2.6.2). Additionally, the difference between the correlation of the CNN and LR models is typically smaller for high wind speed and ice speed, and larger for high sea-ice concentration (Fig. 2.6d–f in section 2.6.1), large distances from the coast, and large depths

(Fig. 2.4 in section 2.6.1). We aim to confirm these findings by comparing the spatial and temporal variability in model correlation (Figs. 2.3a–f & 2.5a–g) to that of the various properties linked to ice motion (Figs. 2.3g–l & 2.5h–j), as well as to the spatial variability of the LR coefficients (Fig. 2.7a–c).

We analyze four spatial divisions (Fig. 2.1a) that are made based on overall model performance and the difference between the performance of the CNN and LR models. The Greenland Sea division (dark blue in Fig. 2.1a) covers the region to the east of Greenland where the model correlation is variable, but the LR largely outperforms the CNN. The Eastern Arctic division (light blue in Fig. 2.1a) represents the region of the eastern Arctic where the correlation is low and the LR outperforms the CNN. The Central Arctic division (light red in Fig. 2.1a) includes the central Arctic, the Beaufort Sea, and the regions to the north of the Canadian Arctic Archipelago. The Baffin Bay division (dark red in Fig. 2.1a) is the region where the model correlation is variable, but the CNN consistently outperforms the LR. The gray shading in Fig. 2.1a indicates regions that are not included in the following analysis. We discuss how the variability in the input parameters is linked to (i) model performance, (ii) the difference between the performance of the CNN and LR models, and (iii) the values for the LR coefficients in each division. We note the distinction between inter-divisional comparisons and analysis within each division, both of which are discussed below.

A summary of the inter-divisional comparisons is shown in Fig 2.9. Here the average values of the metrics and properties are shown for each division, and error bars represent the standard deviation. While the mean over any given division falls within one standard deviation of the mean for the other division for many properties, significance testing shows that for each property the differences between the mean value for each individual division and all other divisions are statistically significant (not shown). Analysis within each division is summarized in Fig. 2.10, which shows the ensemble-averaged correlation between each of the performance metrics and each of the properties related to ice motion within each division. The correlation between the maps of the performance metrics (Figs. 2.3a & f) and the average of the properties

throughout time (Figs. 2.3g–i) are shown in Fig. 2.10a & b. The correlation between the daily time series of the performance metrics and the spatially averaged properties (similar to Figs. 2.5d & g vs. Figs. 2.5h–j, but using daily rather than monthly values) are shown in Fig. 2.10c & d. The properties are compared to the model correlation (circles, Figs. 2.10a & c) and the difference between the CNN and LR correlation, $corr_{CNN} - corr_{LR}$ (triangles, Figs. 2.10b & d). The different divisions are represented by the different colors, as indicated in the legend. Values greater than zero are representative of cases where increases in the property are linked to increases in the model performance metric, while values less than zero indicate an inverse relationship between the property and performance metric.

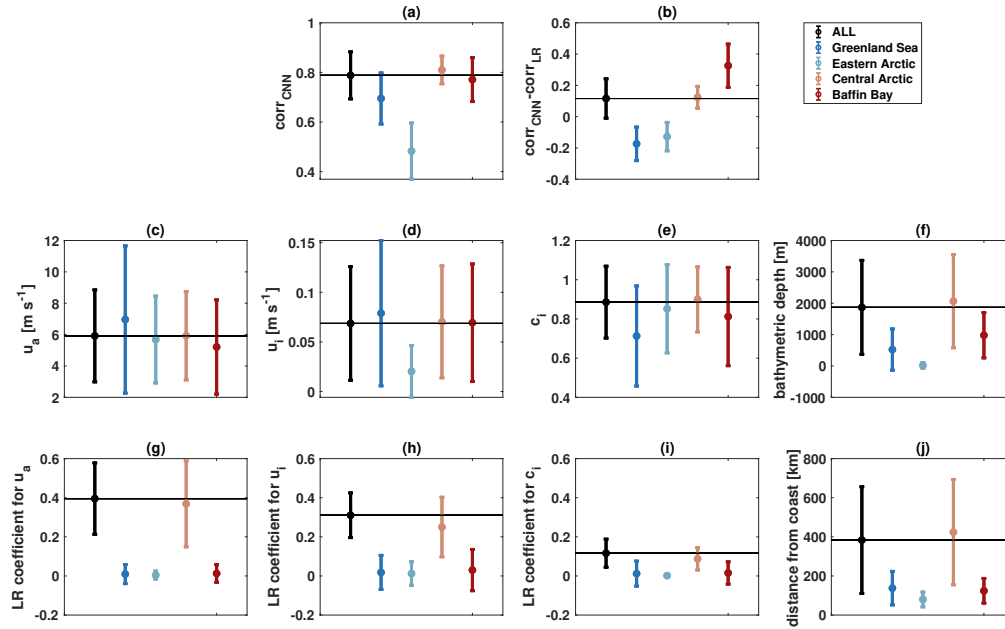


Figure 2.9. Overall mean of the performance metrics, ((a) $corr_{CNN}$, and (b) $corr_{CNN} - corr_{LR}$), and properties related to ice motion ((c) wind velocity, u_a ($m s^{-1}$); (d) ice velocity, u_i ($m s^{-1}$); (e) ice concentration, c_i ; (f) bathymetric depth, d (m); (j) distance from coast, d_c (km); and the LR coefficients for (g) u_a , (h) u_i , and (i) c_i). Different colors represent the different spatial divisions, as indicated in the legend. Error bars represent the standard deviation within each division. The black line in panel each represents the mean value for the overall Arctic ('ALL' in the legend) for comparison.

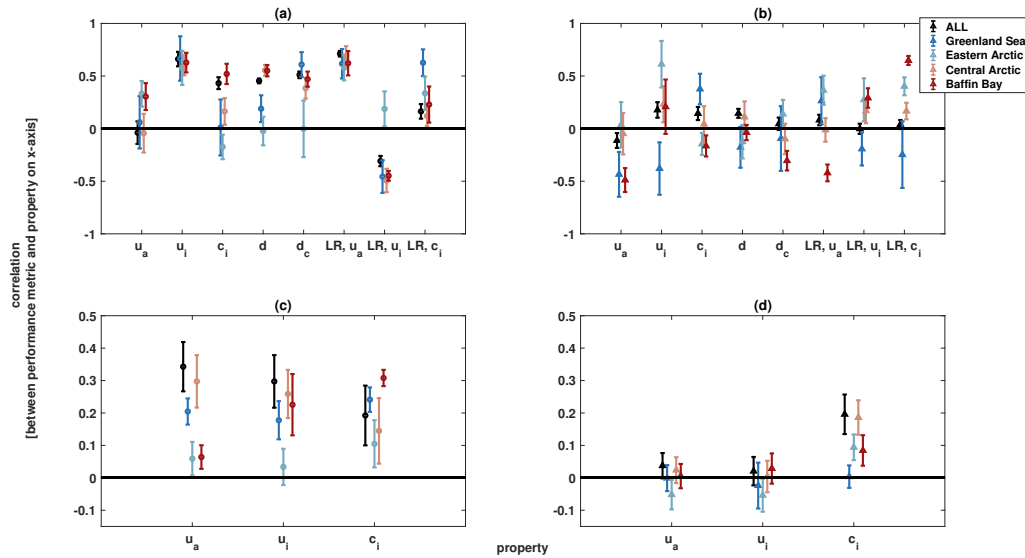


Figure 2.10. Ensemble mean of the correlation between the model performance metrics (circles for $corr_{CNN}$ in a & c; triangles for $corr_{CNN} - corr_{LR}$ in b & d) and the various properties related to ice motion within each of the spatial divisions (different shades of red and blue, as indicated in the legend). Correlations are calculated to understand how (a & b) spatially mapped performance metrics are related to spatial variability in time-averaged wind speed, u_a ; ice speed, u_i ; ice concentration, c_i ; depth, d ; distance from coast, d_c ; LR parameter for wind, LRu_a ; LR parameter for ice speed, LRu_i ; and LR parameter for ice concentration, LRc_i ; and (c & d) temporal variability in performance is linked to daily averages of u_a , u_i , and c_i within each division. Error bars represent the standard deviation of the ensemble runs within each division.

Model predictive skill vs. properties related to ice motion

Inter-divisional comparisons suggest that low correlation of the CNN is typically linked to low depth, distance from coast, and ice speed, which is consistent with results from Fig. 2.6. For example, the Eastern Arctic division has the lowest $corr_{CNN}$, as well as the lowest mean of the properties listed above in comparison to the other divisions (Fig 2.9).

Visual inspection of spatial (Fig. 2.3) and temporal (Fig. 2.5) results also support this. For example, the low $corr_{CNN}$ found in the Eastern Arctic division (Fig. 2.3c) is coincident with low values for depth, distance from coast (Fig. 2.1b–c), wind speed, ice speed, and ice concentration (Figs. 2.3g–i). Temporally, the exceptionally low correlation in the Eastern Arctic

division from November–May (Fig. 2.5b–d) is coincident with values of u_i for the Eastern Arctic division that are lower than all of the other divisions (Fig 2.5i). Additionally, the Central Arctic division exhibits a higher correlation than the other divisions, particularly during May–October, where the Central Arctic has higher u_a , u_i and c_i in comparison to the other divisions. Temporal analysis also shows that divisions that have a lower seasonal minimum c_i also exhibit a lower correlation relative to the other divisions, and in August–September the ordering for both c_i and $corr_{CNN}$ between divisions is: Baffin Bay < Eastern Arctic < Greenland Sea < Central Arctic.

Within each division, large $corr_{CNN}$ is typically related to high depth, distance from coast, wind speed, ice speed, and ice concentration, which is consistent with results from Fig. 2.6. This can be seen in Fig. 2.10a & c, where data points for all divisions are typically greater than zero (above the black line), which indicates that spatial (Fig. 2.10a) and temporal (Fig. 2.10c) variability of the properties listed on the x-axis are linked to variability in the correlation of the CNN. There are a few exceptions to this relationship when comparing spatial variability of performance metrics to the mean field of the properties: large wind speed is linked to low $corr_{CNN}$ within the Central Arctic and the overall Arctic; within the Eastern Arctic division large ice concentration, depth, and distance from coast are linked to low $corr_{CNN}$. Interestingly, many of these exceptions lie within the Eastern Arctic division, where overall depth, distance from coast, wind speed, ice speed, and ice concentration are significantly lower than other divisions. However, the values of these exceptions are within one standard deviation of zero, which indicates neither a positive or negative correlation between the model performance and the respective property. We note that the spatial comparisons (Fig. 2.10a) make use of the mean fields of u_a , u_i , and c_i , while temporal analyses (Fig. 2.10c) look at the daily time series that are averaged over the spatial domain of each division. We use spatial and temporal analyses here as a confirmation of results in Fig. 2.6, but do not expect perfect adherence due to the differences caused by averaging across space and time.

While Fig. 2.10a & c provides a quantitative analysis of the comparisons of spatial (Fig. 2.3c vs. g–i) and temporal (Fig. 2.5d vs. h–j) variability in the model correlation with

respect to these properties, we can also see the link through visual inspection. For example, spatial patterns of high correlation within the Greenland Sea division (i.e. increasing from west to east; Fig. 2.3a–c) are coincident with high depth, distance from coast, ice speed, and wind speed, while low correlation is seen in locations with high ice concentration. Within the Eastern Arctic division, low correlation is largely linked to low depth and ice speed (Fig. 2.10a). High correlation within the Central Arctic division is generally coincident with high depth, distance from coast, ice speed, and ice concentration. Slightly lower correlations are seen in regions with lower values of u_a and u_i (western side), and lower c_i (eastern side and near the Bering Strait). Interestingly, the Beaufort Sea has high values of skill and correlation despite its proximity to land. However, the Beaufort Sea is relatively deep and has exceptionally high mean u_a and u_i in comparison to other coastal regions, properties that are linked to higher model performances (Figs. 2.4d & 2.6a–b). Lastly, high model correlation in the Baffin Bay division (Fig. 2.3) is aligned with large depth (Fig. 2.1b), u_a , u_i , and c_i (Fig. 2.3g–i). These spatial patterns of correlation within each of the divisions tend to be consistent with results from Figs. 2.4d & 2.6a–b, the main exception being for the link between low correlation and high c_i within the Greenland Sea division, and the high correlation found close to the coast in the Beaufort Sea.

Temporally, the seasonal cycle for correlation follows that of u_a , u_i , and c_i , with minimum model correlations occurring during the months of minimum u_a , u_i , and c_i (August–September) for most models and divisions. The exceptions here are the Eastern and Central Arctic divisions where the correlation does not follow the seasonal cycle for c_i . This is likely a result of the low u_a and u_i in the Eastern and Central Arctic division during this time. Additionally, low seasonal variability in correlation within the Central Arctic division could be linked to the relatively small seasonal variations in u_a , u_i , and c_i in comparison to the other divisions.

Difference between predictive skill of the CNN and LR models vs. properties related to ice motion

Inter-divisional analysis suggests that low values for the difference metric, $corr_{CNN} - corr_{LR}$ (the Greenland Sea and Eastern Arctic divisions in Fig 2.9b), are linked to low depth, distance from coast, and ice concentration (the Greenland Sea and Eastern Arctic division in Fig 2.9 e–f & j). Additionally, a low difference metric is linked to high u_a and u_i in the Greenland Sea division. Conversely, low u_i is linked to a low difference metric in the Eastern Arctic division (Fig 2.9d). The high difference metric in the Baffin Bay division is also linked to a lower mean u_a . As noted above, while the mean value of a particular division may fall within one standard deviation of that for other divisions, significance testing shows that the differences between means among divisions for a given property are statistically significant. For the case of c_i , these inter-divisional comparisons are consistent with results from Fig. 2.6d–f, where a high difference metric is linked to high c_i . Additionally, these results are consistent with the relationship between high $corr_{CNN} - corr_{LR}$ and low wind and ice speeds found in Fig. 2.6d–f for the Greenland Sea (u_a & u_i) and Baffin Bay (u_a) divisions, but not the Eastern Arctic division (u_a & u_i).

Visual inspection of spatial (Fig. 2.3f vs. Figs. 2.1b–c & 2.3h–j) and temporal (Fig. 2.5g vs. Fig. 2.5h–j) results also supports this. Spatially, the low difference metric, $corr_{CNN} - corr_{LR}$, in combination with relatively low depth, distance from coast, ice concentration, and exceptionally high wind and ice speeds in the Greenland Sea division compared to the rest of the Arctic is consistent with results in Fig. 2.6d–f. Additionally, temporal analysis shows the difference metric for the Greenland Sea division remains lower than that for the entire Arctic (dark blue line is below black line), while u_a and u_i are higher in the Greenland Sea division than other divisions during the months of October–April. Similarly, for the Eastern Arctic division a relatively low depth, distance from coast, and ice concentration are linked to a low difference metric. However, contrary to patterns found in Fig. 2.6, the difference metric in the Eastern Arctic division is low, while u_a and u_i are also low in both spatial and temporal analyses. The difference metric for the Eastern Arctic division is lower than that for the Greenland Sea division from

January–April, despite lower c_i and higher u_a and u_i in the Greenland Sea division, all of which are expected to contribute to a lower difference metric (Fig. 2.6). Spatially, the high difference metric in the Central Arctic division is linked to high c_i , low u_a , and high u_i relative to other divisions, which is consistent with results in Fig. 2.6, with the exception of the tendency of u_i . However, in temporal analysis of the Central Arctic division, the difference metric is particularly high compared to other divisions when u_i is lower in January–May, which is consistent with results in Fig. 2.6. The notably high difference metric in the Baffin Bay division compared to other divisions is linked to low u_a in both spatial and temporal (December–June in Fig. 2.5g & h) analyses.

Within each division, comparisons of $corr_{CNN} - corr_{LR}$ with the properties related to ice motion are more nuanced, as data points in Fig 2.10b & d do not consistently lie above or below zero for a given property, particularly with spatial comparisons using the mean fields (Fig 2.10b). From results in Fig. 2.6, we would expect points in Fig 2.10b to be above zero for c_i and below zero for u_a and u_i (i.e. high $corr_{CNN} - corr_{LR}$ is linked to high c_i , low u_a , and low u_i), which is only the case for some divisions. The region containing the entire Arctic (black) is consistent with this pattern for all variables on the x -axis, except for u_i . Additionally, these results are consistent with Fig. 2.6 for the following cases: the coincidence of high $corr_{CNN} - corr_{LR}$ with low u_a , low u_i , and high c_i in the Greenland Sea division; high $corr_{CNN} - corr_{LR}$ coincident with low u_a , but high c_i and depth in the Central Arctic region; the coincidence of high $corr_{CNN} - corr_{LR}$ with low u_a in the Baffin Bay region. We find the following exceptions to the trends in Fig. 2.6: the coincidence of high $corr_{CNN} - corr_{LR}$ and low d and d_c in the Greenland Sea division; high $corr_{CNN} - corr_{LR}$ coincident with high u_i and low c_i in the Eastern Arctic division; and high $corr_{CNN} - corr_{LR}$ coincident with high u_i , low c_i , low depth, and low distance from coast in the Baffin Bay division.

Comparisons between temporal variability of the difference metric and the various properties are more straightforward, and tend to show results that are consistent with what is found in Fig. 2.6, where a high difference metric is linked to low u_a , low u_i , and high c_i . This is

true (i.e. data for u_a and u_i exist below the black line, and points for c_i are above), except for in the case of the region containing the entire Arctic, the Central Arctic division, and the Baffin Bay division for both u_a and u_i , as well as the Greenland Sea division for c_i . Additionally, while the ensemble mean value of the correlation between u_i and the difference metric is negative for the Greenland Sea division, it lies within one standard deviation of zero.

Looking at the time series (Fig. 2.5g–j) it is clear that the low difference metric in the Eastern Arctic division from December–May is linked to low u_i and high c_i , which is the opposite of what is expected from Fig. 2.6. Within the Central Arctic division low $corr_{CNN} - corr_{LR}$ is linked to low c_i in June–October, while a slightly higher $corr_{CNN} - corr_{LR}$ from December–May is linked to high u_a and low u_i . Within the Baffin Bay division low $corr_{CNN} - corr_{LR}$ is linked to low c_i (Figs. 2.5g & j & 2.10d): during months of low c_i , the difference metric is not statistically different from zero (May–November, except August), while for all other months the opposite is true, and $corr_{LR} < corr_{CNN}$. Additionally, high $corr_{CNN} - corr_{LR}$ during January–April is coincident with a low u_a . Temporal results from Fig. 2.5g–j tend to be consistent with results from Fig. 2.6, with the following exceptions: coincidence of low $corr_{LR} < corr_{CNN}$ with low u_i and high c_i from December–May within the Eastern Arctic division; coincidence of high $corr_{LR} < corr_{CNN}$ and high u_a from December–May in the Central Arctic division.

Impact of LR parameters on model performance metrics

We find that the performance metrics ($corr_{CNN}$ and $corr_{CNN} - corr_{LR}$) are related to the values of the LR coefficients for the different input parameters (Fig. 2.8 in section 2.6.2). These results come from comparing the LR coefficient at each location (Fig. 2.7a–c) with the mapped values for the performance metrics (Fig. 2.3c & f). We use divisional analyses to confirm the maximum LR coefficient in each division (Fig. 2.9g–i vs. Fig. 2.7f), as well as the relationship between the performance metrics and the LR coefficients within each division (Fig. 2.10a & b vs. Fig. 2.8). We also aim to understand whether the variable with the highest LR coefficient has the strongest relationship to model performance.

Inter-divisional comparisons (Fig. 2.9g–i) show that the mean LR coefficient for u_a is higher than all other coefficients in the Central Arctic division and the region covering the entire Arctic. For all other divisions the mean of the LR coefficients are within one standard deviation of each other and the maximum coefficient within each division is not conclusive. The mean LR coefficient within the overall Arctic and the Central Arctic division (Fig. 2.9g–i) is consistent with what is seen spatially (Fig. 2.7f). We find that variability in model performance is not necessarily linked most strongly to the property that exhibits the dominant LR coefficient within each division (i.e. a high LR coefficient for u_a does not necessarily mean that the correlation between either performance metric and u_a will be stronger than that between the performance metric and u_i or c_i). In other words, the high value of the LR coefficient for u_a in comparison to that for u_i or c_i for the Central Arctic division in Fig.2.9g is not linked to the correlation between model performance and u_a being higher than that for u_i or c_i in Fig 2.10.

In Fig. 2.8, high model correlation is found in locations with large LR coefficient for u_a and c_i , but a low LR coefficient for u_i . Analysis of the LR coefficient within each division (Fig. 2.10a & b) confirms this and shows that high $corr_{CNN}$ is related to high a high LR coefficient for u_a and c_i within all four divisions. The relationship between high $corr_{CNN}$ and a low LR coefficient for u_i is also seen for all divisions except the Eastern Arctic division (light blue in Fig. 2.10a). While the general trend in Fig. 2.8 suggests high correlation to be linked to a low LR coefficient for u_i , it is clear that when $corr_{CNN} < 0.6$ (which is the case for the Eastern Arctic division, where the mean $corr_{CNN}$ is 0.5 ± 0.02), a high LR coefficient for u_i is linked to higher $corr_{CNN}$.

The relationship between a high difference metric and a high LR coefficient for u_a seen in Fig. 2.8 is confirmed within all divisions, except for the Baffin Bay division (Fig. 2.10b), however the Central Arctic division is within one standard deviation of zero. The relationship between high $corr_{CNN} - corr_{LR}$ and a low LR coefficient for c_i seen in Fig. 2.8c is only found within the Greenland Sea division. While the general pattern in Fig. 2.8c suggests a link between high $corr_{CNN} - corr_{LR}$ and a low LR coefficient for c_i , this is largely true where the LR coefficient

for c_i is high (> 0.6), which is the case for the Greenland Sea division (0.69 ± 0.34). When the LR coefficient for $c_i < 0.6$ the opposite is true, and high $corr_{CNN} - corr_{LR}$ is linked to a high LR coefficient for c_i , which is the case for the Greenland Sea, Eastern Arctic, and Baffin Bay divisions. Thus, Fig. 2.10b confirms results from Fig. 2.8.

2.7 Conclusions

2.7.1 A CNN can make skillful predictions of sea-ice motion on one-day time scales.

As sea-ice in the Arctic declines and opens new pathways for maritime transportation, the skill of sea-ice motion predictions becomes increasingly important (Bennett et al., 2020; Cao et al., 2022). This work uses machine learning models to make one-day predictions of sea-ice motion for operational forecasting. We show that a CNN can make skillful predictions of sea-ice velocity and outperforms other statistical models in most instances. In comparison to the other models, the CNN has the benefit of incorporating non-linearities between inputs and spatial information when making predictions. We also show that a grid-wise linear regression (LR) model performs almost as well as a CNN in most instances, and comes with the benefit of decreased complexity in comparison to neural networks. Both the CNN and LR models outperform the baseline PS model. Additionally, we find that the CNN shows improved performance in comparison to the models of Maeda et al. (2020); Kimura and Wakatsuchi (2000) discussed in section 2.3: the correlation of the CNN is as low as 0.4 in the Eastern Arctic, and 0.7 in the Canadian Arctic Archipelago (Fig. 2.3c), where Maeda et al. (2020) find correlation between ice motion and geostrophic wind as low as 0 and 0.4 in the same regions. Lastly, while comparing the model performance to that of a dynamical model was outside the scope of this study, our model was an extension of that presented by Zhai and Bitz (2021) (differences between the two models are identified in Table B.1), which was found to have higher correlations for sea-ice velocity with satellite observations than the CICE5 dynamical model for sea ice.

2.7.2 Model predictive skill and discrepancies between model performances are linked to various properties related to sea-ice motion.

Model performances vary spatially and seasonally, and are linked to variability in properties related to sea-ice motion. Although there are exceptions that come with having different combinations of these properties, in general, better model performance is linked to:

- increased bathymetric depth and distance from the coast
- larger mean values of u_a , u_i , and c_i
- larger LR coefficients for u_a and c_i ; smaller LR coefficient for u_i

The CNN outperforms the LR in most cases. We have shown that the following are related to increases in the performance of the CNN over the LR:

- larger distance from coast and greater bathymetric depth
- smaller mean u_a and u_i , and larger mean c_i
- larger LR coefficient for u_a , and smaller LR coefficients for c_i

Interestingly, the LR model tends to outperform the CNN model in some coastal regions where non-linear effects might be expected to play a large role. However, the locations where this happens exhibit shallow depths, and when coastal waters are deep (i.e. the Beaufort Sea) the CNN outperforms the LR. We note that sharp discontinuities between ocean and land pixels may reduce the quality of the CNN predictions due to the way the CNN incorporates spatial filters and non-local information in its predictions Sonnewald et al. (2021). This may also impact our result that the LR outperforms the CNN at shallower depths because depth increases with increasing distance from the coast. To address this, future analyses we will apply a non-local LR at each grid-point for a more direct comparison between LR and CNN models. However, even with non-localities built in, the LR doesn't apply spatial filters in the same way that the

CNN does, so we may not be able to reproduce the same decreases in performance inherent to the CNN in coastal regions.

The LR typically outperforms the CNN in regions where wind speed is not the dominant LR coefficient: ice velocity is the dominant LR coefficient in the coastal regions of the eastern Arctic, and sea-ice concentration dominates the LR predictions in the coastal region to the east of Greenland. Conversely, wind speed is found to be the dominant LR coefficient wherever the CNN outperforms the LR. This suggests that the relationship between wind velocity and ice velocity includes non-linearities that are captured by the CNN (and not the LR), leading to an improved performance.

We find that larger LR coefficients for a given parameter are not necessarily linked to larger parameter values (e.g. in the Greenland Sea division, ice concentration is the dominant predictor in regions where wind and ice speed are exceptionally high). However, we find that the LR coefficient for wind speed tends to be lower in regions with low mean c_i . This contradicts previous findings, where areas with high c_i are known to exhibit larger internal ice stresses, which leads to a reduction in the dependence of ice motion on wind (Kimura and Wakatsuchi, 2000; Maeda et al., 2020). We note that this particular conclusion does not take into account instantaneous effects, as it is a comparison between a mean c_i over time and a LR coefficient that is descriptive of ice motion over the duration of the study. Future work could decrease the time period over which LR is run to obtain equations that are more descriptive of instantaneous effects such as that of ice stresses due to high c_i . Lastly, we find that variability in model performance is not necessarily linked to the dominant LR coefficient within each region.

2.7.3 Wind velocity plays the largest role in predicting ice velocity.

We find that the spatial average of the wind factor over the Arctic is $0.72\% \pm 0.31\%$ (Fig. B.2). The wind factor is higher for regions in the Central Arctic in comparison to coastal regions, confirming historical results (Thorndike and Colony, 1982; Serreze et al., 1989; Kimura and Wakatsuchi, 2000; Maeda et al., 2020). We also show an average turning angle to the wind

of $24.9^\circ \pm 11.3^\circ$, which is consistent with the cited historical results. Analysis of LR parameters shows that of all of the input predictors, wind velocity has the largest importance in predicting sea-ice velocity. This relationship is particularly strong in the central Arctic, and is reduced in coastal regions. Furthermore, an increased dependence of the models on wind speed is related to increased model performance for the CNN, which provides further evidence as to why the models are not as skillful at predicting ice speed in coastal regions (i.e. ice speed is not as dependent on the training information in these regions). Future work will build off of these results and look at using outputs from machine learning models to understand how the relationship between wind and ice velocity is changing in time as the ice melts.

2.8 Acknowledgements

LH, MRM, and PH were supported by ONR (grant N00014-20-1-2772). MRM was supported by NSF (award OPP-1936222). STG was supported by NSF (award OPP-1936222) and by U.S. Department of Energy (DOE) (Award DE-SC002007). DG was supported by NSF Award 1928305. CMB was supported by NASA (award 80NSSC21K0745). Figures in this report were prepared using MATLAB, Matplotlib: A 2DGraphics Environment Hunter (2007). Colormaps were obtained using the cmocean package (Thyng et al., 2016) and the CubeHelix Colormap Generator (Stephen23, 2023). Special thanks to the coauthors who provided invaluable insight for this chapter: Matthew Mazloff, Sarah Gille, Donata Giglio, Cecilia Bits, Patrick Heimbach, and Kayli Matsuyoshi. We thank our reviewers for their helpful feedback.

Chapter 2 is a reprint of the material as it may appear in Artificial Intelligence for Earth Systems, 2023. We cite this work as: Hoffman, L., Mazloff, M. R., Gille, S.T., Giglio, D., Bitz, C. M., Heimbach, P., and Matsuyoshi, K. (2022). Machine learning for daily forecasts of arctic sea-ice motion: an attribution assessment of model predictive skill. *Artificial Intelligence for Earth Systems*. ©American Meteorological Society. This preliminary version has been accepted for publication in *Artificial Intelligence for the Earth Systems* and may be fully cited. The final

typeset copyedited article will replace the EOR when it is published. Used with permission. The dissertation author was the primary investigator and author of this paper.

2.9 Data Statement

We acknowledge all sources of publicly available data that were used in this study. The JRA55-do dataset can be accessed at <https://climate.mri-jma.go.jp/pub/ocean/JRA55-do/>. Polar Pathfinder Daily 25 km EASE-Grid Sea Ice Motion Vectors, Version 4 are made available by the National Snow and Ice Data Center (NSIDC) and can be accessed at <https://nsidc.org/data/nsidc-0116/versions/4>. Sea Ice Concentrations from Nimbus-7 Passive Microwave Data, Version 1 are made available by the NSIDC and can be accessed at <https://doi.org/10.5067/8GQ8LZQVL0VL>. The International Bathymetric Chart of the Arctic Ocean (IBCAO) are available at https://www.gebco.net/data_and_products/gridded_bathymetry_data/arctic_ocean/. All of the data and files used for processing for this paper can be accessed at <https://doi.org/10.6075/J0X06774>.

Chapter 3

Evaluating the robustness of explainable machine learning (XML) methods for application in regression predictions of Arctic sea-ice motion.

3.1 Summary

Recent advances in explainable ML (XML) methods have shown promise for understanding predictions made by machine learning models. In this study we evaluate explainability methods applied to linear regression (LR) and convolutional neural network (CNN) models that are built to predict present-day ice velocity from inputs of present-day wind velocity, previous-day ice velocity, and previous-day ice concentration. We aim to understand the importance of each of the inputs in predicting sea-ice velocity on one-day timescales. This work focuses on methodology in order to gain a comprehensive understanding of the various XML methods. We compare more established techniques with those that have been more recently introduced to the field. In particular, we compare covariance maps, parameters from linear regression equations, permutation feature importance (PFI) maps using CNN, perturbation experiments using CNN, and layerwise relevance propagation (LRP) methods. This is the first known application of a domain-integrated implementation of the LRP method for a regression problem in the geosciences. We refer to “domain-integrated” or “regional” explainability methods when describing

methods that interpret the importance of the inputs over the entire spatial domain of interest, rather than just at one grid location. The LRP method was developed for ML models built for classification. Here we apply XML to regression rather than classification-type predictions, and thus seek to evaluate the robustness of LRP for its application to regression, rather than classification-type predictive ML models. Outputs from different explainability methods are generally consistent. Results from local methods are consistent in showing the relevance of each of the inputs has a similar radius of influence for the perturbation and LRP methods. However, local methods are inconsistent in determining the relevant importance of the inputs: perturbation indicates wind speed is the most relevant predictor of ice motion, while LRP shows that ice concentration is the most relevant. Domain - integrated methods show similar inconsistencies in the relative importance of the inputs for predicting ice motion: while both the LR parameters and PFI methods show that wind speed is the most important predictor, domain - integrated LRP indicates that ice concentration has a higher relevance in predicting ice motion. We also discuss nuances found in the spatial distribution of relevance between the various domain-integrated explainability methods.

3.2 Introduction

Machine learning (ML) models can make skillful predictions of sea-ice motion on one-day time scales (Zhai and Bitz, 2021; Hoffman et al., 2023). Sea-ice velocity has historically been modeled linearly as a function of wind speed, moving at some fraction of the wind speed (known as the wind factor), and at some angle to the wind (the wind angle; Thorndike and Colony, 1982). On short timescales, the wind alone can account for 70% of the variability in sea-ice motion in the central Arctic region. In regions within 400 km of the coast, the wind-only approximation of ice motion becomes less reliable, as internal stress gradients within the ice become more important (Hibler, 1979; Thorndike and Colony, 1982; Serreze et al., 1989; Kimura and Wakatsuchi, 2000; Maeda et al., 2020). This historical result was seen in the decrease in

performance in coastal regions of linear regression (LR) and convolutional neural network (CNN) models analyzed in Hoffman et al. (2023). These models were designed to predict present-day ice velocity from present-day wind velocity, previous-day ice velocity and previous-day ice concentration. Even with additional information about the ice state (i.e. the previous-day velocity and concentration), the performance of these models dropped in coastal regions. The previous chapter discussed potential mechanisms to explain the decrease in model skill in coastal region (see section 2.7).

We move forward with analysis of the LR and CNN models for predicting ice motion. In the this chapter and the next, we aim to understand how the relationships between ice motion and the model input predictors are changing in response to the changing state of the Arctic. To accomplish this, we employ explainable ML (XML) methods to probe into the black-box of the ML models and gain insight into how each of the inputs is important in making one-day predictions of ice motion. This chapter begins as an analysis of several different XML methods, and the final chapter investigates trends in the relevance of each of the inputs in predicting ice motion.

Investigations in the interpretability of ML predictions have become more common with an increasing reliance on complex neural networks (Linardatos et al., 2021). CNNs are black-box models that incorporate non-linear interactions between the predictors, and it is difficult to interpret the reasoning behind any given prediction. However, several explainability methods have been developed with the aim of untangling the black box nature of the CNN. Several different ways of categorizing these methods have been identified (Haar et al., 2023; Letzgus et al., 2021; McGovern et al., 2019). Here, we analyze perturbation-based and back-propagation-based methods. We further classify methods into those producing local or domain - integrated interpretations of the model predictions (Linardatos et al., 2021).

In this study we analyze the outputs from different XML methods applied to a CNN (a non-linear, black-box model) in comparison to a simplified analysis of the coefficients in a LR model (a linear, glass-box model). We use the coefficients of the LR equation as an

indicator of how important each of the inputs was for the model in predicting the output. We analyze the perturbation and back propagation-based methods applied to a CNN for local and domain-integrated cases. These include local implementations of perturbation and LRP, and domain-integrated implementations of PFI and LRP. We evaluate the outputs from these different XML methods applied to a CNN and compare to the LR case for a robust analysis and to bolster trustworthiness in the interpretations of our models.

3.3 Data

We train linear regression (LR) and convolutional neural network (CNN) models to make predictions of present-day zonal and meridional sea-ice velocity ($u_{i,t}$ & $v_{i,t}$) from inputs of:

- present-day zonal & meridional wind velocity ($u_{a,t}$ & $v_{a,t}$),
- previous-day zonal & meridional sea-ice velocity ($u_{i,t-1}$ & $v_{i,t-1}$), and
- previous-day sea-ice concentration (c_{t-1}).

Daily values of wind velocity, ice velocity, and ice concentration from 1989–2021 are taken from the same satellite and reanalysis sources discussed in chapter 2.4, which can be referenced for further information on data sources and processing steps. In summary, wind velocity is taken from the Japanese Meteorological Agency 55-year atmospheric surface dataset for driving ocean-sea ice models (JRA55-do; Tsujino et al., 2018); ice velocity is from the Polar Pathfinder Daily Sea Ice Motion Vectors, Version 4 (Tschudi et al., 2019); ice concentration is from the Nimbus-7 SMMR and SMSP SSM/I-SSMIS Passive Microwave Data (Cavalieri et al., 1996).

Similarly to chapter 2.4, we re-grid wind and ice concentration to the 25 km EASE-Grid of the Polar Pathfinder Ice Motion product for consistency. We also standardize all data to zero mean and one standard deviation, based on global statistics over the entire temporal domain. Sea-ice velocity and concentration are set to zero rather than ‘NaN’ in regions without ice

because the CNN requires inputs to be numerical values. The Nimbus-7 sea ice concentration data set is missing data in a circular sector centered over the North Pole as a result of the orbit inclination of the satellite. In the analysis in chapter 2 these grid locations were simply filled in with zeros before implementation in model training. For XML studies here, we have filled in this polar hole using imputation methods, replacing locations with ‘NaN’ values with the mean of the 40 nearest neighboring grid points. We opted to apply imputation here because while the polar hole did not significantly impact the overall performance of the model (which was the focus of chapter 2), it did impact the XML outputs that were analyzed in this chapter (not shown).

3.4 Methods

3.4.1 Model Setup

The model frameworks are taken from Hoffman et al. (2023) and further described here.

Linear Regression (LR)

The LR is expressed as equation (3.1):

$$u_{i,t}^* = Au_{a,t}^* + Bu_{i,t-1}^* + Cc_{i,t-1}^* + D. \quad (3.1)$$

Here the inputs and coefficients (i.e. $u_{a,t}^*$ & A , etc.) are complex numbers, with the real and imaginary parts representing the zonal and meridional components and their respective parameters. The LR is applied grid-wise and each grid location has a unique set of LR coefficients. We apply a time-variable mask that only uses grid points and times where the ice concentration is greater than zero. Similarly to Hoffman et al. (2023), we apply ridge regression with a ridge parameter of $\lambda = 10^{-2}$ to prevent unrealistically large LR parameters (Marquardt and Snee, 1975).

We run LR for three different temporal subsets: overall, seasonal, and monthly. For the overall runs, data are split into train and test sets, which include 30 and 2 years of data,

respectively. Ten different LR models are trained, where the years used in the train and test sets are varied; we refer to each of these models as an “ensemble” run. For the monthly runs, we build a different LR model for each month from 1989–2021. For the seasonal runs, a different LR model is built for each season (i.e. January–March, JFM; April–June, AMJ; July–September, JAS; and October–December, OND) from 1989–2021.

Convolutional Neural Network (CNN)

The CNN architecture is illustrated in Figure 3.1. The model is set up with five repeating units of Conv2D-ReLU-MaxPool, followed by a 20% dropout, a flattening, and a dense layer that applies a regression to predict the output at each grid point. The CNN is implemented in python using the Tensorflow/Keras library (Abadi et al., 2015). The parameters are described in Table 3.1 and further details can be found in section 2.5.2. Data are split into train, validation, and test sets with an 88%–6%–6% split. We train ten different CNN models with varying date ranges used the train, validation, and test sets. Similarly to the LR, we refer to each of these models as an “ensemble” run. We use the CNN to apply several different types of sensitivity analyses to understand the importance of the input parameters in prediction sea-ice motion. These analyses include perturbation, permutation feature importance (PFI), and layerwise relevance propagation (LRP), which are discussed in more detail in the following section.

3.4.2 Sensitivity Analyses

We aim to explore the importance of each of the inputs to the ML models (i.e. present-day wind velocity, u_a ; previous-day ice velocity, u_i ; and previous-day ice concentration, c_i) in predicting the output. We apply different types of sensitivity analyses to LR and CNN models, including perturbation and back-propagation-based methods. These can be separated into two types: local and domain-integrated. Localized analyses show the sensitivity of each of the inputs at individual grid locations, whereas domain-integrated analyses provide insight into the overall relevance of each input in predicting the output over the entire analysis domain.

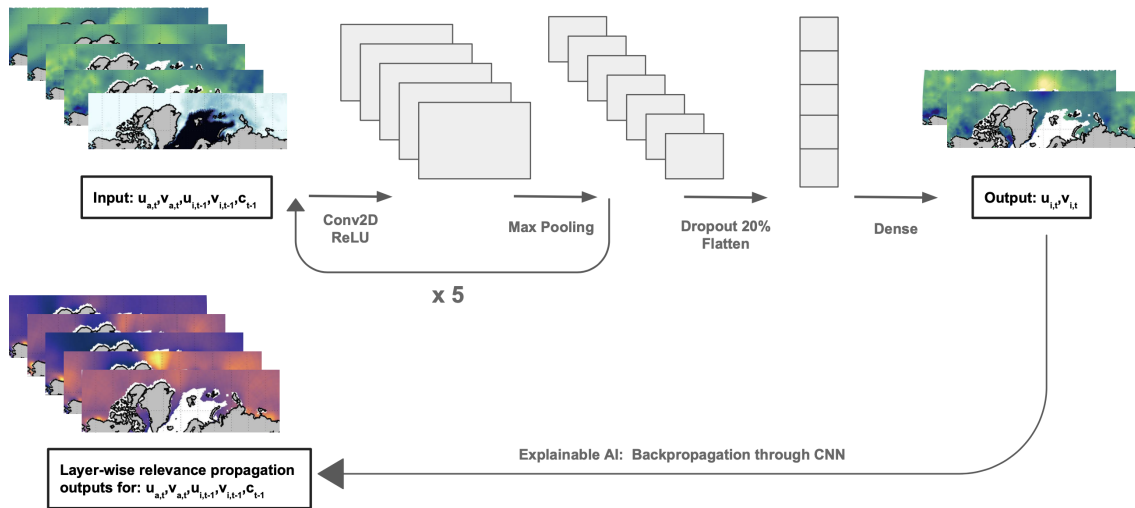


Figure 3.1. Schematic of the convolutional neural network (CNN) used in this study for predicting present-day sea-ice velocity components (outputs) from present-day wind velocity, previous-day sea-ice velocity, and previous-day sea-ice concentration (inputs). This CNN has five repeating units of a 2D convolution with a ReLU activation and max pooling, followed by a 20% dropout layer, flattening, and a dense layer. An example of the heat maps produced using layerwise relevance propagation (LRP) is also shown here. LRP highlights regions of high relevance for each of the inputs in making predictions of sea-ice motion.

Localized methods include analyzing covariance maps of the inputs, perturbation analyses, and gridwise-LRP. Domain-integrated methods include analysis of the LR parameters, permutation feature importance (PFI) applied to the CNN, and a global implementation of the LRP. Each of these methods are described below. We compare the outputs from various sensitivity methods to validate the use of newer methods (i.e. LRP for regression problems) against the more established case (i.e. analysis of LR parameters).

Covariance Maps of Inputs

We evaluate spatial covariance maps of each model input at 17 different locations throughout the Arctic. We note that the covariance is not a metric of sensitivity magnitude for the relationship between ice motion and the input parameters, but use it here to analyze the spatial extent to which each input co-varies in time. This provides a visual of the spatial decorrelation length scale for wind velocity, ice velocity and ice concentration. The spatial covariance pertains

to the linear relationship between two points for a given predictor. The magnitude of the spatial covariance of a particular predictor indicates the degree to which knowledge at one spatial location permits at least partial knowledge of information at another location (Wunsch, 2006). This metric gives information about the distance over which the 17 chosen analysis points the variations in each of the inputs are synchronous. We analyze spatial covariance over the full temporal duration of the data (i.e. 1989–2021).

Linear Regression Parameters

The linear regression parameters provide information about where each of the inputs is important for predicting sea-ice motion. In comparison to sensitivity studies using the CNN, LR only shows information about linear relationships. We build different LR models over varying temporal subsets of data, as discussed in section 3.4.1. These include implementations for the overall (trained on 30 years of data), seasonal (a separate model trained on each season from 1989–2021), and monthly (a separate model trained on each month from 1989–2021) cases. Each grid point has a unique set of LR parameters relating the inputs to the output. Analysis of the LR parameters from the overall case allows comparison with the sensitivity studies applied to the CNN that have domain-integrated results; we use this method as a baseline for comparison with explainability methods applied to a more complex model (i.e. a neural network). We take the average of each of the LR parameters over the ten ensemble runs to create a map indicating the overall impact of each input at each location.

Perturbation Analysis

Perturbation analyses provide information about the sensitivity of a ML model prediction to a perturbed element (McGovern et al., 2019; Ivanovs et al., 2021). The perturbed element is considered important if the model prediction changes as a result of the perturbation. We are interested in understanding the relative importance of each of the input predictors at various geographical locations. Thus, we apply perturbation analyses for each input at 17 different

locations throughout the Arctic. We follow a procedure similar to that of Sinha and Abernathey (2021). After the CNN has been trained we apply the model to the standard test data set (2020–2021) to make a prediction. Next, a new test data set is made for each of the five inputs and each of the 17 hand-picked locations, for a total of $17 \times 5 = 85$ perturbation runs. Here, the input is perturbed at the chosen location while keeping the rest of the variables and locations fixed. For each perturbation, we subtract a fraction (1/2) of the standard deviation of the input. This statistic is taken globally over all space and time. The model is run on each perturbed test data set to make a prediction. For each perturbation, the root mean square difference (RMSD) is calculated between the prediction made using the perturbed and the non-perturbed input. We evaluate the spatial extent of the RMSD by summing the square of the difference in the temporal dimension for each grid point (equation 3.2).

$$RMSD = \sqrt{\sum_i^n \frac{(\hat{y}_i - y_i)^2}{n}}, \quad (3.2)$$

Permutation Feature Importance (PFI)

Another way to evaluate the sensitivity of the ML model to the various inputs is through permutation feature importance (PFI) methods (Radivojac et al., 2004). Similarly to perturbation, in PFI the relative importance of the inputs is determined by the extent to which the ML model predictions are impacted by changing elements of the inputs. To apply PFI, we randomly shuffle the spatial values of each input at each time step for the test data set (i.e. 2020–2021). We then use the shuffled data to make a prediction with a trained CNN. The performance (correlation and skill, equations 2.4 & 2.5) of the CNN is calculated and compared to a standard case without randomization. The magnitude of the change in performance is a metric for how important the randomized input is in predicting the output. We analyze this importance for each input spatially at each grid point. From this analysis we can determine the input with the overall highest importance, in addition to spatial locations that are identified as important by the CNN. This PFI

method has been applied to ML predictions made in geosciences in the form of classification problems involving sea ice (Shen et al., 2017) and weather event detection (Molina et al., 2021; Lakshmanan et al., 2015), as well as for regression problems for predicting motions of wave gliders (Amador et al., 2021).

Layerwise Relevance Propagation (LRP)

We use a method called layerwise relevance propagation (LRP) to trace the explanation of CNN predictions (Bach et al., 2015; Montavon et al., 2015, 2019, 2018). LRP informs the “relevance” of individual input parameters by back-propagation through the neural network to the first layer. The process proceeds as follows: (i) the CNN is run, and weights and biases are frozen; (ii) a prediction output is propagated backwards onto the frozen weights and biases of the different layers of the CNN and feature relevance is learned through a process governed by conservation, where neurons that contribute the most to the higher-layer receive most relevance from it; (iii) the output is a heat map showing the “relevance” of each input parameter at each mapped location (latitude & longitude). This process is repeated for each prediction output at each location, producing a daily relevance heat map for each input time step and grid point. We proceed with analysis of LRP in two different ways: locally and domain-integrated. For the local case, we pick 17 different latitude-longitude locations in the Arctic and analyze the relevance heat maps produced from each of these locations in comparison to other explainability methods that are confined to one grid point (i.e. covariance maps and perturbation). For domain-integrated analyses, we create integrated heat maps by running LRP for 219 selected points throughout the Arctic (Fig. 3.2) and averaging the relevance maps produced by each of these locations at each time step. This allows for comparison of LRP to more spatially comprehensive explainability methods, such as the LR parameters or PFI outputs.

We run the LRP analysis on the testing data for consistency with other XML analyses. Here we use the iNNvestigate package (Alber et al., 2019) with the ‘sequential preset A’ configuration of LRP that applies different rules at different layers of the model, where $LRP-\alpha_1 \beta_0$ is applied

for convolutional layers and LRP- ϵ is applied for dropout, flatten, and dense layers. This custom configuration was chosen based on suggestions from Montavon et al. (2019).

While LRP methods were developed for classification rather than regression (Bach et al., 2015; Montavon et al., 2019), there have been examples of LRP being applied for regression in other fields (Dobrescu et al., 2019; Rahman et al., 2021; Schnake et al., 2020). Additionally, Letzgus et al. (2021) discuss methods to extend LRP to regression problems. These methods involve re-training the CNN with respect to a carefully chosen reference value (i.e. subtracting the reference value from the outputs before training). For consistency in comparison with other explainability methods, we show LRP relevance outputs where the reference value is not applied for local and domain-integrated implementations in this chapter. We discuss applications of the reference value to future work in the final chapter.

In comparison to other explainable AI techniques (i.e. perturbation and gradient-based), backward propagation techniques, such as LRP, have the advantages of producing explanations from a single forward / backward pass (Letzgus et al., 2021). To our knowledge, this is the first application of using a spatially-integrated LRP for a regression problem in the geosciences. For more details, an overview of using LRP for applications in geosciences is provided by Toms et al. (2020).

3.5 Results

3.5.1 Localized Sensitivity Studies

We compare the sensitivity studies applied at specific grid points: covariance, perturbation and localized-LRP. We run each of these analyses for 17 different locations (red points in Fig. 3.3g). While we have noted that the covariance is not the same as a sensitivity between the model inputs and ice motion, we refer to the “relevance” of each of the inputs within each of these methods to simplify the discussion. The sensitivity outputs are normalized by dividing by the top 0.5% maximum value for each particular method to create similar scales for comparison among




Figure 3.2. Locations used in integrated analysis for global LRP implementation (blue dots). Regions of the Arctic are labeled for reference during discussion.

the different methods. Thus, when we refer to “relevance” we are referring to the covariance, the RMSD, and the LRP output, each normalized to the 0.5 % maximum value for each of the covariance, perturbation, and LRP methods, respectively. We compare relevance maps that indicate the importance of each of the inputs (e.g. wind speed) in predicting the output (i.e. sea-ice velocity). The spatial extent of the relevance metric for each sensitivity method is shown for location 9 (Fig. 3.4); an extensive look at this set of maps for each location can be found in the appendix (Fig. C.1 – C.17). The rows represent the different sensitivity methods, and the columns are the inputs.

Spatial covariance decreases with an increasing distance from the analysis point for all

predictors. Sea-ice concentration has the highest and most far-reaching spatial covariance, with locations throughout the Arctic having highly correlated values for ice concentration (Fig. 3.4c). For both wind and sea-ice velocity, spatial covariance remains high within a radius of the analysis location, but drops off significantly with increasing distance (Fig. 3.4a & b). Within that radius, wind and ice velocity have a similar spatial covariance, whereas the covariance of ice velocity is slightly larger than that for wind at greater distances. While we show results for location 9, the same patterns tend to hold for all 17 analysis points. We note that when analyzing the correlation of for the case where the seasonal cycle has been removed (not shown), the radius of influence of the spatial covariance for ice concentration is lower, while that for wind and ice speed does not change much. This suggests that the seasonal cycle dominates the covariance of ice concentration, but not wind or ice speed. Thus removal of the seasonal cycle could play a role in the outputs of explainability methods, particularly for the case of which input has the largest relevance in predicting the output.

Perturbation analyses on the CNN show decreasing importance with an increasing distance from the analysis point for all predictors (Fig. 3.4d–f). Wind velocity has the highest importance and radius of influence, followed by ice concentration and then ice velocity. The normalized perturbation relevance score (i.e. the RMSD in equation 3.2 normalized to the 0.5 % maximum value) for wind speed remains above 0.3 even for locations across the Arctic (i.e. off the east coast of Greenland), while that for ice speed drops down to zero at the same location. In comparison to all other methods and inputs, the perturbation relevance score for wind remains above 0.8 for the largest distance away from the analysis point.

Similarly, LRP applied at the analysis location shows high localized importance that decreases with increasing distance from the analysis point for all predictors (Fig. 3.4g–i). Here, ice concentration shows the highest relevance in predicting ice velocity, followed by wind velocity and the previous-day ice velocity. LRP relevance tends to be more non-localized in comparison to perturbation analyses, with locations of high relevance farther away from the analysis point. This is particularly true for ice concentration, which was also shown to have a

larger decorrelation length scale in the covariance maps (Fig. 3.4c). Interestingly, some of the locations far from the analysis point where LRP shows high relevance for each of the inputs are the same locations with high spatial covariance for sea-ice concentration; this occurs, for example, in the Eastern Arctic and to the east of Greenland. Additionally, while areas of high relevance in perturbation tend to extend coherently within a certain radius of the analysis point, the spatial extension of high LRP relevance is not coherent and instead high relevance is also found at locations farther away (i.e. Fig. 3.4d vs. g). Overall, the LRP relevance seems to have more similarities in spatial variability to the covariance maps than does the perturbation relevance. This difference in the spatial patterns of relevance suggests that LRP and perturbation pick up on areas of high relevance through different mechanisms.

For a comprehensive analysis of the relevance at each of the 17 locations (labeled in Fig. 3.3d), we show the mean and standard deviation relevance for each of the inputs and each analysis method (Fig. 3.3a–c). Variations in the relative importance of the inputs discussed for location 9 (Fig. 3.4d–f) tend to hold for all locations. Sea-ice concentration has the largest spatial covariance and LRP relevance in comparison to the other inputs (i.e. the red line is above the orange and yellow lines in Fig. 3.3a & c), while wind velocity tends to have the largest relevance in the perturbation analyses (i.e. yellow line is on top in Fig. 3.3b). Additionally, if a location tends to have a higher relevance for one parameter, it has high relevance for the others as well. For example, the LRP relevances of the wind, ice speed, and ice concentration parameters are higher for locations 4, 7, and 9 than the other locations (Fig. 3.3c).

We also analyze the extent to which the relevance values vary with distance from the analysis location for all locations and each of the sensitivity methods and inputs (Fig. 3.5). Here, we only show data within 2000 km of the analysis point because our interest concerns the regions of high relevance. The red lines represent exponential fits, and the legend shows the r^2 value and the e-folding distance, which is a measure of the radius of influence of each method. Relevance decreases exponentially with increasing distance from the analysis point. This is true for all sensitivity methods and inputs. However, the covariance of the ice concentration remains high at

large distances from the analysis point, as indicated by a relatively low r^2 for the exponential fit, and an e-folding distance that is one order of magnitude (OM) higher than that for all other methods and inputs (Fig. 3.5c). For all methods, the ‘relevance’ score for ice concentration has the highest radius of influence (Fig. 3.5c, f, and i). We note that the LRP has more locations with higher relevance at larger distances from the analysis point than perturbation for the case of sea-ice concentration (Fig. 3.5i vs. f); the LRP also has areas with relevance scores above 0.2 at distances greater than 2000 km, which are not shown here. Results in Fig. 3.5 include all locations. We also calculate the radius of influence for each location individually and show the mean and standard deviation over all locations for each input and method (Fig. 3.5j–l). The mean radius of influence for the LRP and perturbation methods are similar, and fall within one standard deviation of each other for each of the inputs. Generally, the radius of influence for LRP and perturbation are similar for each location, with exceptions at a few locations (i.e. L2, L11, L12). Additionally, between the inputs, the mean e-folding distances are within one standard deviation for each method, except for the case of the spatial covariance for ice concentration.

3.5.2 Domain-Integrated Sensitivity Studies

We compare sensitivity studies that analyze domain-integrated relevance of each input in predicting ice motion: linear regression parameters, PFI, and domain-integrated LRP. The sensitivity outputs are normalized by dividing by the top 0.5% maximum value for each particular method to create similar scales for comparison among the three methods. We compare relevance maps that indicate the importance of each of the inputs to the ML models in predicting sea-ice motion and analyze the spatial extent of the relevance produced from each of these methods in Figs. 3.6. Based on historical results (Thorndike and Colony, 1982) we expect wind velocity to have the highest relevance out of all of the inputs.

The LR parameters are calculated on data from 1989–2019, where data from 2020–2021 were saved for use as the ‘test’ data set. The CNN used for both PFI and LRP was trained on data from 1989–2017, where data from 2018–2019 were used for validation and 2020–2021 for

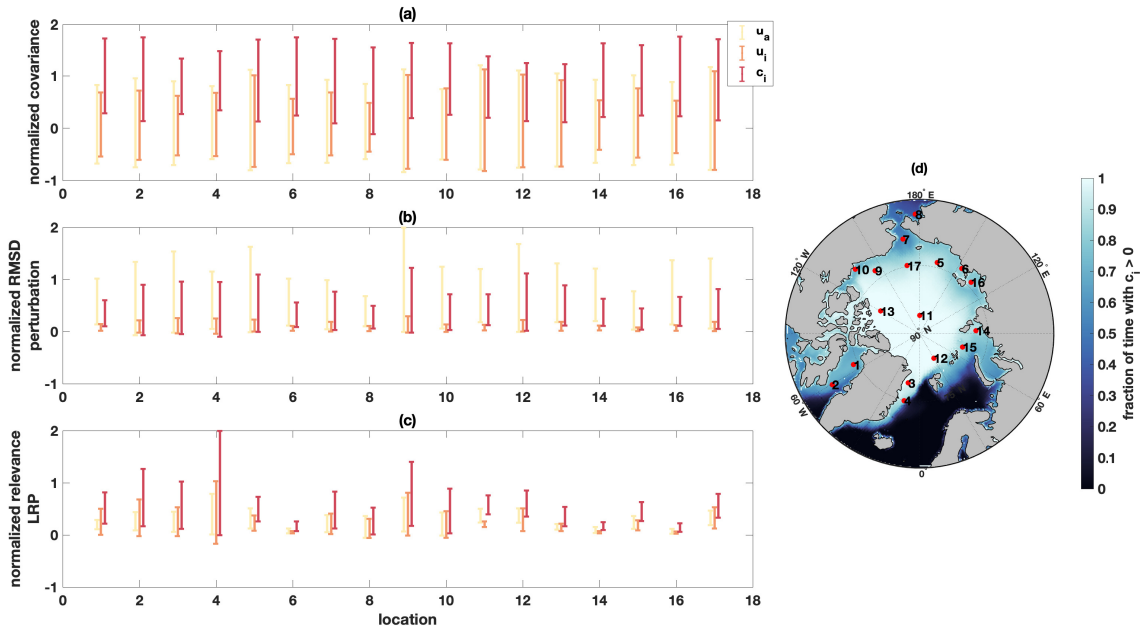


Figure 3.3. Spatial mean and standard deviation of the “relevance” scores for the covariance, perturbation, and LRP sensitivity methods: (a) spatial covariance; (b) normalized RMSE from perturbation analysis; (c) normalized relevance score from LRP for each of the inputs (yellow, wind speed, u_a ; orange, ice speed, u_i ; red, ice concentration, c_i) at each location (x-axis). For consistency in the magnitude of the axes among the methods, each subplot is normalized by dividing by the maximum spatial mean and standard deviation of the ‘relevance’ value for each method. (d) Map showing the fraction of time the ice concentration is greater than zero, as well as the location of each point of analysis (red dots; labeled 1–17).

testing. Output from the PFI and LRP methods is shown here for the test years (2020–2021). We note that for a more comprehensive analysis in the future we will analyze these XML methods for ten different model runs where the years for the train, validation, and test data sets are shuffled as to account for temporal trends.

Spatial analysis of the LR parameters indicates that wind velocity has the highest relevance in predicting ice motion for the LR model, particularly in the central Arctic (dark blue in Fig. 3.6d). Previous-day ice velocity is the second largest LR parameter in the central Arctic, and the highest in coastal regions (Fig. 3.6f & h).

The PFI method also shows that wind velocity is the most important predictor. However, the spatial variability in relevance for the PFI is not consistent with the LR case. For PFI, the wind parameter maintains high relevance throughout the Arctic, even (and especially in some

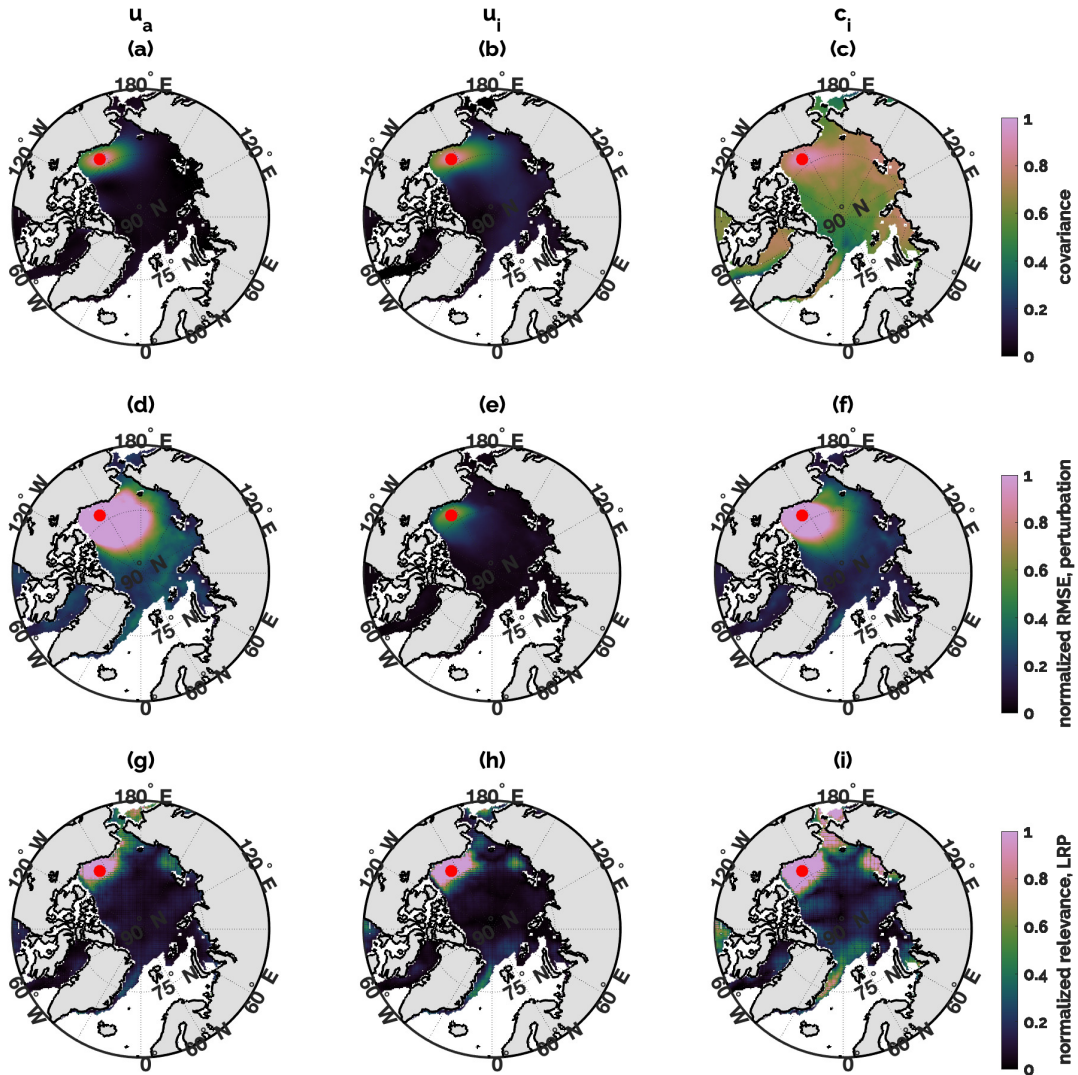


Figure 3.4. Results from localized sensitivity studies for each of the inputs at location 9, indicated by the red dot. The columns represent each of the different inputs: (a), (d), and (g), wind speed, u_a ; (b), (e), and (h), ice speed, u_i ; and (c), (f), and (i), ice concentration, c_i). The rows represent the different sensitivity methods: (a–c) Spatial covariance; (d–f) normalized RMSE from perturbation analysis; (g–i) normalized relevance score from LRP. Each row is normalized by dividing by the maximum ‘relevance’ value of the spatial mean for each method.

cases, such as in the Eastern Arctic) in coastal regions. Spatial patterns in the relevance of the ice speed for the PFI analysis show similarities to that of the LR: the relevance of ice concentration is higher for coastal regions than the central Arctic (i.e. the Canadian Arctic Archipelago, Bering Strait and Eastern Arctic), and in particular for the region to the east of Greenland. There are also

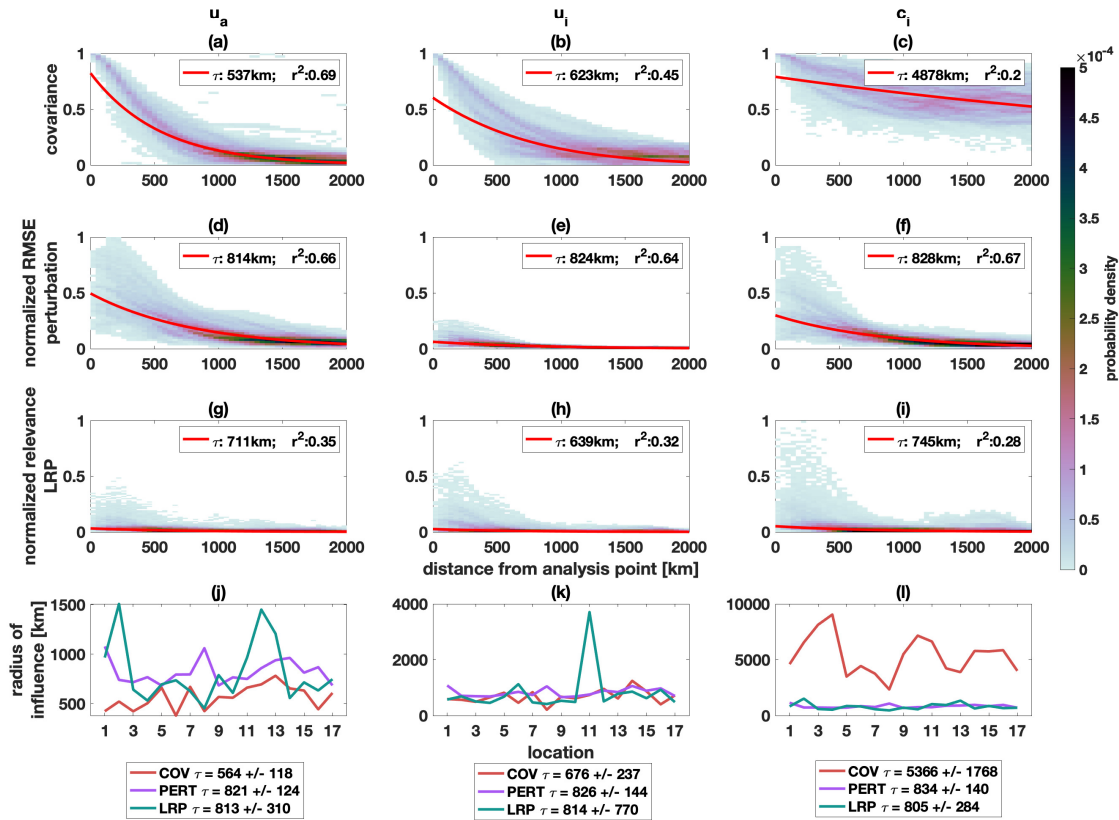


Figure 3.5. (a–i) Probability density of the relevance of each localized sensitivity study as a function of the distance from the point of analysis for each of the inputs and for all locations. The columns represent each of the different inputs: (a), (d), and (g), wind speed, u_a ; (b), (e), and (h), ice speed, u_i ; and (c), (f), and (i), ice concentration, c_i). The rows represent the different sensitivity methods: (a–c) normalized spatial covariance; (d–f) normalized RMSE from perturbation analysis; (g–i) normalized relevance score from LRP. Each row is normalized by dividing by the maximum ‘relevance’ value for each method. The red lines represent exponential fits to the data. The legend gives the e-folding distance for that fit and gives a measure of the radius of influence for each of the relevance methods. The legend also shows the r^2 values for each fit; and $r^2 > 0.12$ is statistically significant with 95% confidence based on the degrees of freedom for each fit. (j–l) Radius of influence (i.e. e-folding distance) for each location individually for each input (columns) and relevance method (colored lines: red for covariance (COV), purple for perturbation (PERT), and green for LRP). The mean and standard deviation over all locations for each input and method are shown in the legend, located below each panel.

similarities between the spatial patterns in the relevance of ice concentration between the LR and PFI methods, where ice concentration has a higher relevance in Baffin Bay for both methods.

The domain-integrated LRP implementation shows that sea-ice concentration is the most relevant predictor of ice velocity (yellow in Fig. 3.6i). Based on the other XML methods we

do not expect this, and instead expect wind velocity to have the highest relevance. We further discuss these results in section 3.6. While the magnitude of relevance for LRP varies between the inputs, the spatial patterns between the inputs are all similar and tend to show higher relevance in the coastal regions (Fig. 3.6i–k; particularly off the coast of Greenland, and in the Beaufort, Bering, and East Siberian Seas). For sea-ice concentration, spatial patterns of relevance for LRP are similar to the LR parameters in that they show a relatively high relevance off the coast of Greenland (Fig. 3.6k vs. g). For ice velocity, similar spatial patterns exist between the standard LRP, the PFI, and the LR parameters in that higher relevance values are found in the coastal regions (Fig. 3.6j vs. b vs. f). The spatial patterns in relevance for wind velocity are different for all three methods. For example, the standard LRP shows relatively high relevance in coastal regions, particularly in the Beaufort Sea and East Siberian Sea; these are the regions where PFI conversely has the lowest relevance for wind velocity, and that the LR parameters also show low importance in the wind velocity (Fig. 3.6i vs. e vs. a).

3.6 Discussion and Conclusions

We analyze various XML methods for interpreting and understanding predictions made by machine learning models. Among these are analyzing the spatial covariance of the model inputs, perturbation analysis, localized-LRP, analyzing the coefficients in the LR equation, PFI, and a domain-integrated implementation of LRP. We separate into localized and domain-integrated analyses, and within each category, we compare the methods. This is the first known application of a domain-integrated implementation of the LRP method for sea ice. In this study, we confirm the feasibility of this method through comparison with other more conventional forms of XML (i.e. analysing LR parameters).

Locally, perturbation methods suggest that wind velocity is the most relevant input for predicting ice motion, followed by ice concentration and previous-day sea-ice velocity. Conversely, LRP shows that sea-ice concentration is the most relevant predictor, followed by

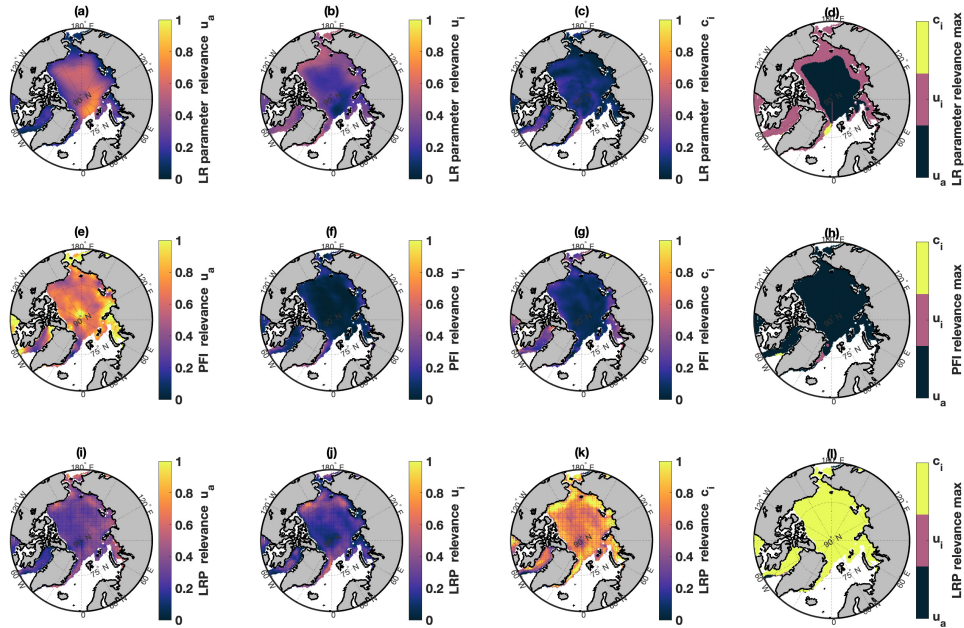


Figure 3.6. Results from global sensitivity studies for each of the inputs. The columns represent each of the different inputs: (a), (e), and, (i), wind speed, u_a ; (b), (f), and (j), ice speed, u_i ; (c), (g), and (k), ice concentration, c_i ; and (d), (h), and (l), the maximum parameter at each location. The rows represent the different sensitivity methods: (a–d) LR parameters; (e–h) PFI; (i–l) LRP. We normalize by dividing by the top 0.5% maximum relevance value of the spatial mean for each method.

wind and then previous-day ice velocity. Relevance decreases exponentially with increasing distance from the analysis point for all methods and inputs. Local implementations of LRP exhibit similar radii of influence to perturbation methods, which fall within one standard deviation of each other for each of the inputs. One difference between methods is that relevance scores greater than 0.2 exist at distances greater than 2000 km from the analysis point for the LRP only (not shown). This suggests that the LRP may be able to capture the non-local information that the CNN is using more effectively than the perturbation method.

For domain-integrated analyses, analysis of LR parameters and PFI similarly show that wind velocity is the most relevant input for predicting ice velocity. However, there are differences in the spatial variability of the relevance heat maps between the two methods. For example, LR parameters show that the relevance of wind speed decreases in coastal regions, where previous-

day sea-ice velocity becomes the most important predictor. On the other hand, PFI shows that wind velocity is consistently the most relevant predictor throughout the Arctic.

We find that outputs from LRP are similar to other sensitivity methods for both local and global cases. However, there are a few contradictions, particularly when it comes to the relative importance of the various inputs in predicting sea-ice velocity. Local and global implementations of LRP show that sea-ice concentration has the highest relevance for predicting sea-ice velocity, while all other methods indicate that wind is the input with the highest importance. Based on the physics of sea-ice motion, we would expect wind to have the largest relevance if the behaviour was linear. This has been the case historically, and wind has been shown to explain up to 70% of the variability in ice motion on short time scales (Thorndike and Colony, 1982).

Based on these expectations, we suspect nuances in the LRP or non-linearities in the CNN that require further analysis. First, we note that for local analyses, the spatial variability of LRP relevance more closely follows that of the covariance than does the perturbation relevance, especially for the case of sea-ice concentration. We noted in section 3.5.1 that the models were run for the case where the seasonal cycle was not removed, and that the spatial variability in the covariance of sea-ice concentration was sensitive to the removal of the seasonal cycle. In future work we will run the CNN and LRP with the seasonal cycle for sea-ice concentration removed to see if this has an impact on the relevance heat maps for sea-ice concentration.

Additionally, from the global analyses we noted that the LRP was highly sensitive to the format of the inputs (not shown). In section 3.4.1 we discussed the fact that LRP can be run to produce relevance heat maps with units (i.e. m/s of output explained per m/s of input). To achieve this, the CNN needs to be run with the inputs standardized to zero mean and one standard deviation, but with the output retaining its original units. After the model is trained, the inputs are un-standardized and returned to their original units before they are analyzed using LRP. The LRP relevance heat maps discussed in this study are for the fully standardized case. In future work we will evaluate the difference in LRP outputs between the standardized (no units) and non-standardized (relevance has units) cases. In future work we aim to understand the nuances

in the impacts of standardization before implementing LRP, and think this will provide clarity on why the LRP shows differences in the relative importance of the various inputs in comparison to other XML methods.

Table 3.1. Details about training data for the CNN and LR models, and model architecture (i.e. parameters and hyperparameters) for the CNN.

Training Data	Inputs	Present-day wind velocity from JRA55-do, $u_{a,t}$ & $v_{a,t}$ Previous-day sea-ice velocity from Polar Pathfinder Sea Ice Motion, Version 4, $u_{i,t-1}$ & $v_{i,t-1}$ Previous-day sea-ice concentration from Nimbus-7 Passive Microwave, Version 1, c_{t-1}
	Outputs	Present-day sea-ice velocity from Polar Pathfinder Sea Ice Motion, Version 4, $u_{i,t}$ & $v_{i,t}$
	Input Size	361 x 361 x 5
	Output Size	361 x 361 x 2
Model Architecture	Layer 1	Conv2D with ReLU: filter size = (3,3), stride = (1,1); Max Pool: filter size = (2,2), stride = (2,2); number of filters = 7
	Layer 2	Conv2D with ReLU: filter size = (3,3), stride = (1,1); Max Pool: filter size = (2,2), stride = (2,2); number of filters = 14
	Layer 3	Conv2D with ReLU: filter size = (3,3), stride = (1,1); Max Pool: filter size = (2,2), stride = (2,2); number of filters = 28
	Layer 4	Conv2D with ReLU: filter size = (3,3), stride = (1,1); Max Pool: filter size = (2,2), stride = (2,2); number of filters = 56
	Layer 5	Conv2D with ReLU: filter size = (2,3), stride = (1,1); Max Pool: filter size = (2,2), stride = (2,2); number of filters = 112
	Layer 6	Dropout, 20%
	Layer 7	Flatten
	Layer 8	Dense
	Optimizer	Adam
	Activation Function	ReLU (slope coefficient, $\alpha = 0.1$)
	Loss Function	Normalized Root Mean Square Error
	Regularizer	L2 with $\lambda = 0.01$
	Epochs	50
	Batch Size	365
Train, Validation, & Test Years	1989-2017 2018-2019 2020-2021	

Chapter 4

Explainable machine learning (XML) for evaluating trends in the drivers of variability in Arctic sea-ice dynamics.

4.1 Summary

The dynamical properties of sea-ice in the Arctic are changing in response to the diminishing sea-ice cover. In this study we apply Explainable ML (XML) methods to a linear regression (LR) model that is built to predict present-day ice velocity from inputs of present-day wind velocity, previous-day ice velocity, and previous-day ice concentration. We aim to understand the importance of each of the inputs in predicting sea-ice velocity on one-day timescales, and to assess how the relationships between these properties and ice motion are changing in response to the changes in the state of Arctic sea-ice. We put particular emphasis on the relationship between wind speed and ice speed (referred to as the wind factor), and investigate trends in the spatial and temporal variability in the relevance of wind in driving ice motion. Overall we find that the ice is becoming more responsive to wind forcing; however we note variability in these changes seasonally and spatially. We analyze this spatio-temporal variability to understand how changes in the responsiveness of ice motion to the wind are linked to changes in ice concentration. We find that locations that exhibit a relatively low wind factor also have a low ice concentration, and that locations with negative trends in sea-ice concentration tend to exhibit negative trends in the

LR parameter for wind speed. Conversely, the ice is becoming more responsive to wind forcing in locations where changes in ice concentration are not significant. These findings contradict the physical understanding of the relationship between ice motion and wind speed and their dependence on ice concentration. Increases in the wind factor have been historically linked to decreases in ice cover and increases in drift speed. These changes have been explained as decreases in ice concentration leading to a regime of free-drifting that is more responsive to wind forcing (Sprenn et al., 2011; Kwok et al., 2013; Maeda et al., 2020). In conjunction with the expected behavior of the wind factor, our results suggest that remote effects are coming into play. We also discuss other mechanisms that may play a role in our findings in comparison to historical results, including the impacts of geographical features, grid-resolution, and the changing interactions between the ocean-ice-atmosphere systems in response to an overall diminishing sea-ice state. Based on these, we suggest localized, in-situ analyses as a future pathway for investigation to better understand how the relationship between ice drift and wind speed is related to changes in the ice state.

4.2 Introduction

Changes in the dynamical properties of sea ice have been linked to a decline in the sea-ice cover in the Arctic (Zhang et al., 2012). For example, widespread increases in sea-ice drift speed have been attributed to reduction in sea-ice thickness in both models (Tandon et al., 2017; Docquier et al., 2017) and observations (Rampal et al., 2009; Sprenn et al., 2011; Zhang et al., 2012), suggesting a link between sea-ice kinematics and sea-ice decline. This evidence supports a mechanism by which sea-ice thinning leads to a reduction in strength, which reduces internal stresses and allows for more deformation, fracturing, and a faster drift (Tandon et al., 2017). Zhang et al. (2012) used a numerical model to show that changes in the dynamical properties of sea ice are linked to reductions in sea-ice volume, concentration, and thickness. In turn, sea-ice motion also impacts the distribution of sea-ice thickness, which may further impact melting

(Untersteiner et al., 2007). Zhang et al. (2000) identified links between large-scale decadal changes in ice thickness and ice dynamics through mechanisms related to the dynamical impact on advection, growth, and lateral melt. Additionally, winter ice redistribution is an important factor in determining summer ice area, with correlation coefficients between the two found to range from 0.5 to 0.9 in some areas of the Arctic (Kimura et al., 2013). The studies relating drift speed to concentration and thickness have been cited as “necessary, but not sufficient” to fully explain the relationship (Docquier et al., 2017). Understanding the coupling between ice motion and its drivers is important for understanding and predicting the state of sea ice. In this study we address these issues by using machine learning methods to gain insight into the drivers of variability in sea-ice motion.

As changes in the ice progress, there will also be changes in the response of the ice to the dynamic mechanisms that drive its motion, particularly the relationship to winds and ocean currents (Serreze et al., 2007). For example, studies focusing on the wind factor (defined as the ratio between ice speed and wind speed) have shown that the ice is becoming more responsive to wind forcing as it melts (Spren et al., 2011; Maeda et al., 2020). As ice cover becomes thinner, ice motion approaches a state of free drift and ice becomes more responsive to wind forcing (Zhang et al., 2012). Spren et al. (2011) have shown that sea-ice drift speeds in the Arctic have seen an upward trend of $10.6\% \pm 0.9\%$ per decade from 1992-2009, with ranges between -4% and 16% per decade, depending on location. In the Central Arctic, a fraction of this observed trend has been found to be explained by an increase in wind speed, but not over the entire basin. In other coastal locations, Spren et al. (2011) concluded that the trend is likely explained by thinning of the ice cover (thickness and concentration). In another study (Maeda et al., 2020), the wind factor exhibits an increasing linear trend from 2002–2016 in many regions of the Arctic during all seasons. These studies support the notion that as the ice melts it is becoming more responsive to wind forcing. We build upon this analysis and use a multiple linear regression (LR) model to understand how the relationship between ice speed and wind speed is changing in time.

In this study we apply LR models to make one-day predictions of sea-ice motion from

satellite and reanalysis sources. We expand on previous work by emphasizing on interpreting the relationships between the input data and sea ice motion, and understanding how these relationships vary in space and time. We analyze the LR coefficients to characterize the relationship between ice and wind in order to improve understanding of how it is changing with changes in the state of sea-ice. We refer to the LR parameter for wind speed as the ‘wind factor’ throughout the rest of this study. We note that our use of the term ‘wind factor’ is slightly different than historical uses of the term, as we apply it to a multiple LR model, whereas historically the term was used to describe the ratio between ice and wind speed in a simple LR model.

4.3 Data

Data used for this chapter are the same as those from the previous two chapters 2.4 and 3.3. For a complete description refer to chapter 3.3.

4.4 Methods

4.4.1 The Models: Linear Regression (LR) and Convolutional Neural Network (CNN)

We use implementations of the LR model from the previous chapter, and refer to section 3.4.1 for a complete description of model architecture and setup. In short, this model is set up to predict present-day ice velocity from present-day wind velocity, previous-day ice velocity, and previous-day ice concentration.

In the following sections, we analyze outputs from the monthly and seasonal LR models, where there are different LR coefficients at each grid point for each month or season of each year, respectively. We use the coefficients in the LR equation (equation 3.1) over these different timescales to calculate yearly trends in the importance of each of the inputs for predicting ice motion on one-day timescales. We combine the coefficients for the zonal and meridional velocity components to analyze the speed (i.e. the LR parameter for wind speed, $LR, u_a =$

$$\sqrt{(LR, u_a)^2 + (LR, v_a)^2}.$$

4.4.2 Trends in the ‘Relevance’ of the Inputs

We analyze seasonal and monthly implementations of the LR to understand yearly trends in the importance of each of the inputs for predicting ice motion. We compare these to the seasonal and monthly values for the inputs. The LR is applied with a time-varying mask excluding locations where the sea-ice concentration is zero. This impacts the degrees of freedom for each implementation of the LR, which is used to determine the significance of the fit and further discussed below.

For the monthly implementation of LR, we employ separate LR coefficients for each grid location, month, and year. These models are used to analyze yearly trends in the importance of the inputs for each month. We find this trend by first calculating the spatial mean of the LR parameter for each month and year. Values of the LR parameter are only included in the spatial mean if they are statistically significant, i.e. if the r^2 value for the particular equation is greater than the critical r^2 for the degrees of freedom for the given month, year and grid point. We then apply a linear least-squares fit to the spatially averaged LR parameters for each month to obtain a yearly trend in each parameter for each month from 1989–2021. Additionally, we apply least squares fits for the months of each year that exhibit the maximum and minimum mean sea-ice concentration to obtain linear yearly trends for times of maximum and minimum sea-ice concentration.

For the seasonal LR, there are separate LR coefficients for each grid location, season, and year. The seasonal models are used to analyze maps of the yearly trend in the LR coefficients of each of the inputs. For this analysis, we apply a least squares regression to the magnitude of the LR parameter for each grid point, input, and season over the 1989–2021 period to obtain a map of the yearly slope of the parameter throughout the Arctic. Before doing this, we mask the LR parameters for any locations, seasons, and years where the fit was not statistically significant. We also mask out regions where the yearly trend in the LR parameter for each season is not

statistically significant.

4.5 Results

We analyze the seasonality of the LR parameters and compare them to the seasonality in ice concentration to understand the relationship that the LR parameter for each input has with sea-ice concentration. We also analyze trends in the LR parameters over time to understand how the relationship between each of the model inputs and sea-ice motion is changing. We compare these trends to trends in the inputs to the model (i.e. wind speed, u_a ; ice speed, u_i ; and sea-ice concentration, c_i) to understand how changes in the ice state are impacting the drivers of ice motion. Analysis of the yearly trends in these parameters are shown and discussed below in the case of overall trend for each month and spatial trend for each season.

4.5.1 Overall LR Parameters and Trends

The monthly LR parameters for wind speed, LR, u_a , ice speed, LR, u_i , and ice concentration, LR, c_i , are shown from 1989–2021 (Fig. 4.1a–c). We also show the monthly wind speed, ice speed, and ice concentration for this time period (Fig. 4.1d–f). The black lines indicate the monthly mean, and the error bars represent the standard deviation of each parameter calculated over the region in the Arctic where sea-ice concentration is greater than zero for more than 20% of the year (i.e the regions shown in Figs. 3.4 and 3.6). The slopes of the red and blue lines indicate the trend in the monthly mean for months of maximum and minimum sea-ice concentration, respectively. The slope and r^2 values for these trends are indicated in the legend, where an $r^2 > 0.12$ indicates a statistically significant trend at 95% confidence.

We find that the trends in all LR parameters are statistically significant during months of minimum and maximum ice concentration, with the exception of the trend for the LR parameter for ice velocity during months of maximum ice concentration (i.e. red line in Fig. 4.1b). The LR parameters for wind speed, LR, u_a , and ice concentration LR, c_i increase throughout time during both months of maximum and minimum sea-ice concentration (red and blue lines in

Fig. 4.1a & c), with a larger slope during months with minimum ice concentration. The LR parameter for ice speed, LR, u_i also increases throughout time during the months of minimum sea-ice concentration for each year (positive slope in the blue line in Fig. 4.1b), but the slope is much smaller than that for wind speed and ice concentration parameters.

We find that the trends in both sea-ice velocity and sea-ice concentration are statistically significant during months of minimum and maximum sea-ice concentration (red and blue lines in Fig. 4.1e & f), while the trends in wind speed are not significant (red and blue lines in Fig. 4.1d). Ice speed is increasing at a rate of $0.087 \text{ cm s}^{-1} \text{ year}^{-1}$ and $0.057 \text{ cm s}^{-1} \text{ year}^{-1}$ for months of minimum and maximum ice concentration, respectively (Fig. 4.1e). Conversely, sea-ice concentration is decreasing throughout time for months with both maximum and minimum ice concentration, with a larger decreasing trend seen during months of minimum ice concentration (Fig. 4.1f).

From these analyses, increases in monthly mean ice speed and decreases in monthly mean ice concentration are linked to increases in the mean LR parameters for wind speed, ice speed, and ice concentration (Fig. 4.1). The largest slope for the LR parameters is seen for the LR parameter for wind speed (Fig. 4.1d), followed by that for ice concentration and then ice speed. The r^2 values for the LR parameters show a similar comparison during months of minimum ice concentration: largest for the LR parameter for wind speed, followed by that for ice concentration and then ice velocity. Overall, these trends suggest that as ice concentration decreases, ice becomes more responsive to wind forcing, and ice velocity increases. Additionally, predictions of ice motion on one-day timescales are becoming more sensitive to ice concentration (Fig. 4.1f). These trends represent the spatial average of changes in the LR parameters, and in the following analyses we address spatial distributions in these trends.

4.5.2 Spatial Variability of LR Parameters Compared to Inputs

We show results from seasonal implementations of LR to further analyze the relationship between the inputs and the LR parameter for wind speed, i.e. the wind factor. In Figure 4.2,

we show the mean wind factor, ice speed, and ice concentration taken over all years for each season. The columns represent the various seasons (January–March, April–June, July–September, and October–December), and the rows are the wind factor, ice speed, and ice concentration, respectively.

The wind factor is relatively high for the Central Arctic and decreases in coastal regions for all seasons (Fig. 4.2a–d). In all seasons, the lowest value for the wind factor is seen in the Eastern Arctic. The wind factor is also particularly low in the Canadian Arctic Archipelago, with the exception of July–September. Spatially, the highest values for the wind factor occur in the Beaufort Sea and in the region northeast of the Greenland Sea from 0–60°E and 80–90°N. Seasonally, the wind factor is the highest from October–November.

Sea-ice velocity shows spatial patterns similar to the wind factor: lower values in the Canadian Arctic Archipelago and Eastern Arctic, and higher values in the Beaufort Sea and in the region northeast of the Greenland Sea (Fig. 4.2e–h). Seasonally, the highest values of ice speed are also found in October–November. One difference is that ice speed is particularly high in the Greenland Sea for all seasons except July–August, where the wind factor is low for this season and region.

Sea-ice concentration exhibits a seasonal cycle: it is high throughout the Arctic in January–March, experiences decreases in the peripheral seas and reaches a minimum in July–September, with particularly low values in the Eastern Arctic, Bering Sea, Kara Sea, and Baffin Bay (Fig. 4.2i–l). Sea-ice concentration remains high in the Canadian Arctic Archipelago and the Central Arctic throughout the year.

4.5.3 Spatial Trends in LR Parameters and Inputs

The spatial variability of yearly trends in the LR parameter for wind speed is compared to the trend in sea-ice velocity and concentration for each season (Fig. 4.3). In this figure, the columns represent the different seasons (January–March, April–June, July–September, and October–December), and the rows represent the LR parameter for the wind speed, LR, u_a ; the ice

speed, u_i ; and the ice concentration, c_i , respectively. We found trends in the LR parameters for ice speed and ice concentration and trends in the wind speed were not statistically significant, and therefore do not show them here. Gray in Fig. 4.3 represents areas where the given yearly trend is not statistically significant at 95% confidence, as judged by the r^2 value for the trend.

Overall we see largely positive trends in the LR parameter for wind speed, LR, u_a (Fig. 4.3a–d), throughout the Arctic and for all seasons. This trend ranges from 2–8% per year and varies spatially throughout the Arctic, though it is particularly high for the Beaufort Sea and Kara Sea from January–June, and for the Central Arctic and Canadian Arctic Archipelago for all months. The exception here is for the months of July–December, where regions in the eastern Arctic exhibit negative trends in the LR parameter for wind speed (i.e. decreases of around 2–4% per year the region between 70°–85°N latitude and from 30°E–120°W longitude). Additionally, the peripheral and coastal seas also exhibit negative trends in the LR parameter for wind speed from January–June.

Ice drift speed is generally increasing. Ice speed shows a consistent positive trend of about 5% per year throughout all seasons and in most regions of the Arctic (Fig. 4.3e–h). The only regions with negative trends in ice speed are some parts of Baffin Bay in January–March and in July–September, and a small region near the islands in the Laptev Sea for July–September. The largest increases in ice speed are seen in the Beaufort Sea for all seasons.

Ice concentration largely exhibits a negative trend throughout all seasons and most regions of the Arctic (Fig. 4.3i–l). From January–June the trends in ice concentration are generally not significant, with the exception of a small negative trend in ice concentration in some of the peripheral seas and in some coastal regions; these trends are roughly - 1% per year. During the months of July–December there is a negative trend in ice concentration throughout most of the Arctic. The most drastic negative trends are found between 70°–85°N latitude and from 30°E–120°W longitude; this regions includes the Beaufort Sea, East Siberian Sea, Laptev Sea, and Kara Sea. In these regions sea-ice concentration is decreasing at a rate of 2–5% per year.

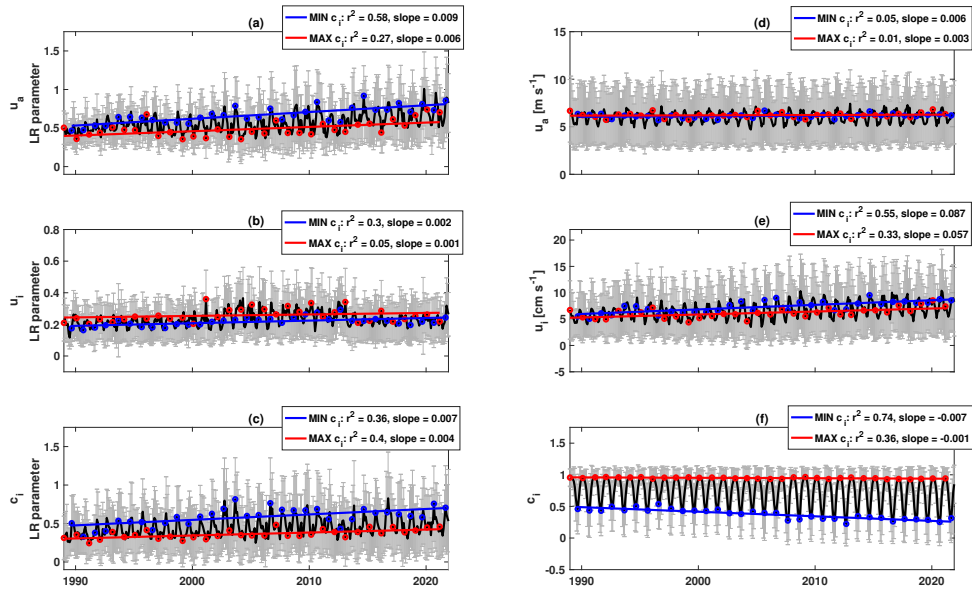


Figure 4.1. Monthly LR parameters and model inputs from 1989–2021. LR parameters for (a) wind speed, LR, u_a , (b) ice speed, LR, u_i , and (c) ice concentration, LR, c_i . Monthly (d) wind speed, u_a , ice speed, u_i , and (f) ice concentration. In (a)–(f), black lines and error bars represent the monthly mean and standard deviation taken over all grid locations. Red and blue lines represent least-squares fits showing the yearly trend in (a)–(f) for months of minimum (blue) and maximum (red) sea-ice concentration in each year. Slopes and r^2 values of these trend lines are indicated in the legend of each panel, where an $r^2 > 0.12$ indicates a statistically significant slope at a 95% confidence level.

4.6 Discussion

4.6.1 How is the wind factor related to ice concentration?

Analysis of the yearly trend in the monthly mean values of each of the model inputs and LR parameters indicates that an increasing LR parameter for wind speed is linked to decreases in ice concentration (Fig. 4.1a & f). This is consistent with the mechanism discussed in section 4.2: a decrease in ice concentration leads to free-drifting ice that is more responsive to wind forcing. We refer to this as the ‘free drift mechanism’ throughout the remainder of this discussion. This phenomena has been supported by several studies that looked at the relationship between the LR parameter for wind speed (i.e. the wind factor) and attributed increases in the wind factor to decreases in ice concentration (Kimura and Wakatsuchi, 2000; Spreen et al., 2011; Zhang et al.,




Figure 4.2. Map of seasonal mean for (a–d) the LR parameters for wind speed, LR, u_a , (e–h) the ice speed, u_i , and (i–l) the ice concentration, c_i from 1989–2021. Each column represents a different season (i.e. January–March, April–June, July–September, October–December).

2012; Kwok et al., 2013; Cole et al., 2017; Maeda et al., 2020). Interestingly, while we find this relationship to hold for the trends in the monthly mean, it is not always the case that the spatial variability in the mean values or trends of ice concentration and the wind factor align in the same way. We continue with a discussion of the nuances in spatial variability.

The Canadian Arctic Archipelago shows consistencies with what is expected, and we find high ice concentration to be linked to low values of the wind factor (Fig. 4.2a–d & i–l). The exception is during July–September, where the wind factor is higher in this region while ice concentration remains high. We note that this is the season with the lowest ice concentration throughout the Arctic, and that the higher wind factor in the region could be attributed to remote effects of lower ice concentration placing the ice in free drift and increasing its responsiveness to wind forcing.

Conversely to what is expected, the Eastern Arctic shows lower values for the wind factor coincident with low ice concentration (Fig. 4.2a–d & i–l). Additionally, lower ice concentration




Figure 4.3. Map of seasonal trend in (a–d) the LR parameters for wind speed, LR, u_a , (e–h) the ice speed, u_i , and (i–l) the ice concentration, c_i from 1989–2021. Each column represents a different season (i.e. January–March, April–June, July–September, October–December). Red and blue regions indicate positive and negative trends, respectively. Grey indicates regions where trends are not statistically significant.

in the Beaufort Sea, Bering Sea, and Baffin Bay are also linked to low values for the wind factor. We note that the wind factor is particularly low in the East Siberian and Laptev Seas for all seasons, even when ice concentration remains high (i.e. January–March). The low wind factor here could be attributed to geographical effects inhibiting the responsiveness of the ice to wind forcing, rather than increases in the ice stresses that result from a higher ice concentration. This would explain the patterns seen in this region for the January–March season.

However, this mechanism (i.e. geographical) does not fully explain the expansion in the area of low wind factor in the the Beaufort-Baffin-Eastern Arctic region in July–September. Based on the free drift mechanism we would expect a higher wind factor in regions of low ice concentration; however we find the opposite. Here, we find a link between low ice concentration and low values of the wind factor, which could be explained by one of the following mechanisms:

(i) grid resolution and (ii) responsiveness of ice to ocean currents. For the grid resolution mechanism, we first note that the ice concentration is a measure of the fraction of the grid box that is covered by ice. Therefore, if floe sizes are small enough, the 25 km resolution of the grid may not be high enough to effectively detect the relationship between wind and the drift of individual floes, particularly if there are multiple floes and a distribution of wind speeds within the grid cell. This could be addressed in future work by applying the same LR methods to the ice drift product from synthetic aperture radar (SAR), which has a spatial resolution of 10 km (Saludo and Hackett, 2020), or by using localised, in-situ data sets. The free drift speed of ice is dependent on the relative contributions of the atmosphere and ocean to ice drift. These are related to the atmosphere and ocean drag coefficients, which depend on (i) frictional force based on the surface roughness of the ice on the ocean and atmosphere side (the skin drag) and (ii) the ice morphology (the form drag) Lepparanta (2011). In this work, we do not include information about ocean currents in our models, but acknowledge they could play a role in the spatio-temporal variability in ice drift. In future work, even something as simple as looking at the effects of the long-term mean ocean circulation could provide useful insight into the relative importance that wind and ocean currents play in driving ice motion.

The highest values of the wind factor are found from October–December in locations where the ice concentration is relatively high (Fig. 4.2d & l). However, this season has a comparatively low overall ice concentration. Here, we suspect that remote effects play a role in the high wind factor through the free drift mechanism: lower overall ice concentration leads to a state of free drift in which ice is more responsive to wind forcing.

Throughout the Arctic and in all seasons we largely see that regions with a high wind factor have a high ice speed (Fig. 4.2a–d & e–h). The exception is the particularly high ice speed coincident with a low wind factor in the Greenland Sea for all seasons except July–August. This region is known as the Fram Strait, and is identified as a region where ice is exported out of the Arctic Ocean. The wind factor is relatively low in this region, but the LR parameter for ice speed (not shown) is high, indicating that ice drift is controlled by persistence, or more likely, the ocean

currents.

We find that spatial analyses show positive trends in the LR parameter for wind speed throughout most regions in the Arctic (Figs. 4.3i–l & 4.3a–d), which is consistent with historical results (Sprenn et al., 2011; Zhang et al., 2012; Kwok et al., 2013; Maeda et al., 2020), and findings in Figure 4.1. However, we find that areas that exhibit the largest negative trend in ice concentration are those that exhibit negative trends in the LR parameter for wind velocity (Figs. 4.3i–l & 4.3a–d). Here we refer to the region of the eastern Arctic (i.e. the region between 70°–85°N latitude and from 30°E–120°W longitude) from July–December, which exhibits some of the largest decreases in ice concentration, coincident with a negative trend in the LR parameter for wind speed. This is inconsistent with the free drift mechanism, which hypothesizes that lower ice concentration leads to a regime of free drift where ice is more responsive to wind forcing. However, the link between decreasing ice concentration and increasing wind factor supported by the free drift mechanism is not negated just because there is not a local correlation between the two. The free drift mechanism could have remote effects. In other words, decreases in ice concentration in one part of the Arctic could free up the ice in other regions to become more responsive to wind forcing. This hypothesized ‘remote effect of free drift’ mechanism is consistent with the overall decrease in sea-ice concentration and maps in Figure 4.3a–d that largely show a positive trend in the LR parameter for wind for most seasons and most regions of the Arctic. This explanation would be consistent with what is known about the physics of ice motion, historical results, and the trends shown in Figure 4.1.

In future work we will use information about the ice thickness and the linear kinematic features (LKFs) in the ice to further investigate this relationship. For example, we may find that ice is not one cohesive structure due to the presence of ridges or leads (i.e. converging or diverging ice floes) in regions where spatial patterns show a link between decreases in ice speed and decreases in the wind factor; this structural nature could impact the motion of the ice and will be further investigated in future work.

We also note that most of the Arctic exhibits positive trends in ice speed (Fig. 4.3e–h),

which further supports the mechanism that decreases in ice concentration lead to a regime where more of the ice is in free drift and able to respond more freely to wind forcing, which increases the ice drift speed (Spreen et al., 2011; Zhang et al., 2012; Kwok et al., 2013; Maeda et al., 2020). Trends in drift speed have been shown to be more substantial in regions with larger decreases in ice concentration (Fig. 4.3e–h & i–l), which has been attributed to decreases in the mechanical strength of ice associated with thinning (Zhang et al., 2012). Further investigations of the relationship between ice drift and wind speed discussed above will be able to provide further evidence to support this mechanism.

4.6.2 Future Work: LRP with CNN

In future work we will replicate the above analysis using the LR model, but instead use a CNN with LRP to understand trends in the relationship between ice drift and wind speed. We will use a CNN with the same architecture discussed in section 3.4.1 (Fig. 3.1). We will employ layerwise relevance propagation (LRP) as an explainability method in order to gain insight into the relevance of each of the inputs in predicting the output. The LRP will be applied using the iNNvestigate package (Alber et al., 2019) with the ‘sequential preset A’ configuration. As discussed in section 3.4.2 we will use a CNN that has been trained against a baseline reference value because we are using LRP for a regression problem (as opposed to classification; Letzgus et al., 2021; Mamalakis et al., 2023). Use of a baseline changes the CNN prediction question from ‘how important are each of the inputs in predicting ice motion different from zero?’ to ‘how important are each of the inputs in predicting ice motion to be different from the baseline?’ (Mamalakis et al., 2023). This is an important distinction, because here we aim to understand how the relationship between ice motion and wind speed is changing as the ice cover decreases. Therefore, the baseline reference value is chosen to be 0.8% of the wind speed and at an angle of 25.3° to the right of the wind, based on the historical linear relationship between ice speed and wind speed (Thorndike and Colony, 1982). The reference value is applied to the output for each time step and grid point.

The LRP can also be applied in a manner that produces relevance scores with units. We will follow steps discussed in section 3.4.2 to train a CNN and run LRP on the outputs. The LRP is applied for the entire duration of the data, and produces a relevance heat map for each input at each day from 1989–2021. Similarly to the LR parameters, we will look at yearly trends in the monthly and seasonal averages of these relevance values to understand how they are changing with time.

The analysis of the CNN will be similar to that done for the LR. However, slight differences exist due to the inherent nature of the ‘relevance’ outputs from these models. For the LR we analyze the coefficients of the equations applied at different timescales (i.e. monthly and seasonally). Applying the LR over shorter periods of time does not impact the overall skill of the model in the same way that it does a neural network, which relies on large amounts of data to make skillful predictions (not shown). We will not perform the CNN over shorter time periods because this does not allow enough data for the model to make skillful predictions. However, we will obtain ‘relevance’ results for the CNN for each day with implementation of the LRP method for explainability. From this point, we will calculate monthly and seasonal mean values to obtain data similar in structure to those of the LR analysis discussed above. We will use the monthly mean LRP relevance value for each input averaged spatially throughout the Arctic to analyze the overall yearly trend in the importance of each input in predicting ice motion for the CNN for each month, and for months of maximum and minimum sea-ice concentration for each given year. For seasonal analyses, we will average the LRP relevance score at each location individually for each season and year, and then calculate the yearly trend in relevance for each season mapped throughout the Arctic. In comparison to the LR, the CNN incorporates information about non-linear and non-local interactions between the inputs, and is therefore more skilled at making predictions of ice motion (Hoffman et al., 2023).

4.7 Conclusions

This study has analyzed how the relationship between ice drift and wind speed (i.e. the wind factor) is related to ice concentration. The primary mechanism by which the wind factor is related to ice concentration is based on what is known about the physics of sea-ice motion and how it is related to wind forcing and internal ice stresses. This ‘free drift mechanism’ links decreases in ice concentration to a regime where the more of the Arctic sea-ice is in free drift, which causes it to be more responsive to wind forcing, which in turn leads to increases in the ice drift speed. Several studies have cited increases in the ice drift speed (Kwok et al., 2013) and increases in the wind factor (Sprenn et al., 2011; Maeda et al., 2020), and related these changes to an overall decrease in ice concentration (Cole et al., 2017). In this study, we highlight ways in which our results are consistent with historical mechanisms, and suggest mechanisms that may be at play for locations where inconsistencies arise.

- The free drift mechanism: Our analyses of the spatial average of the wind factor, ice speed, and ice concentration are consistent with the free drift mechanism. We see positive trends in the yearly wind factor and ice speed coincident with negative trends in sea-ice concentration. Spatial analyses also suggest the free drift mechanism plays an important role in the relationship between the wind factor and ice concentration. For example, in the Canadian Arctic Archipelago regions with high ice concentration exhibit a low wind factor.
- Remote effects of the free drift mechanism: From spatial analyses, we find an overall positive trend in the wind factor throughout the Arctic, along with an overall negative trend in ice concentration. However, these trends do not necessarily coincide spatially. We hypothesize that the free drift mechanism exhibits remote effects where the positive trends in the wind factor result from non-localized decreases in ice concentration that lead to an overall regime of free drift where the ice is more responsive to wind forcing. We suggest

steps that can be taken to test this hypothesis moving forward.

- **Geographical features:** In the Eastern Arctic low values for the wind factor are coincident with low ice concentration, particularly for the months of January–March. This is inconsistent with the free drift mechanism, which suggests a higher wind factor in this region. Therefore, we hypothesize that geographical features (i.e. islands) are creating regions that impede the responsiveness of ice motion to wind in these regions.
- **Grid resolution:** The portion of the Eastern Arctic with a low wind factor expands during months of low ice concentration. This expansion cannot be explained by the geographical features mechanism, and it is inconsistent with the free drift mechanism through which we would expect to see a higher wind factor in conjunction with lower ice concentration. We hypothesize that grid resolution could play a role here. Because the ice concentration is a measure of the fraction of each grid cell that is occupied by ice, grid cells with low ice concentration could be made up of many ice floes that are much smaller than the 25 km grid resolution. If the spatial variability of the wind speed occurs on scales smaller than 25 km, each of these floes could experience a vastly different motion based on spatial variability in the wind. Thus, while the free drift mechanism could still play a role, our ability to see this effect could be inhibited by the spatial resolution of our data.
- **The relative importance of ocean currents versus wind speed in driving local ice motion:** We found that locations with lower ice concentration have a lower wind factor. Additionally, locations with the largest decreases in ice concentration are not coincident with locations with a positive trend in wind factor, and instead exhibit a negative trend in the wind factor. Here, we largely refer to the region in the Eastern Arctic between 70°–85°N latitude and from 30°E–120°W longitude. We note that this study does not include information about the ocean currents, and that the relative role that wind and ocean currents have could be important in determining the role that wind speed plays in driving ice motion.

Chapter 5

Conclusions and Future Work

Understanding the exchanges of freshwater within the ice–ocean–atmosphere system is important because the global freshwater cycle is predicted to amplify in response to expected changes in climate (Yu et al., 2020). Changes in climate are predicted to cause increases the frequency of extreme AR events and seasons over the CCS (Dettinger, 2011; Payne et al., 2020). Additionally, we are experiencing a regime-shift to a “New Arctic” that is characterized by decreases in ice extent and thickness, warming and freshening of the Arctic Ocean, and enhanced ice–ocean–atmosphere coupling, all of which will play a role in the freshwater cycle in the Arctic and around the globe (Solomon et al., 2021). This study assesses freshwater exchanges between ice, ocean, and atmosphere off the coast of California and in the Arctic. We divide the analysis into three main parts: (i) characterizing the impact of precipitation from ARs on upper ocean salinity in the CCS (chapter 1); (ii) assessing the viability of using ML models to predict and understand changes in Arctic sea-ice dynamics (chapters 2 & 3), and (iii) characterizing the changing relationship between sea-ice drift and wind speed (chapter 4).

5.1 Surface salinity response to precipitation from atmospheric rivers in the California Current System.

While studies of the ocean’s response to precipitation have been extensive in the tropics, there remain gaps in the understanding of the role large precipitation events play in modulating

upper ocean salinity in the subtropics off the coast of California. In this study we analyze the impact of precipitation from ARs on the surface ocean salinity in the CCS. The first part of this dissertation addresses the following questions related to ARs in the CCS:

- How does the surface salinity respond to precipitation from ARs in the CCS on seasonal and event timescales?
- How does atmospheric forcing in the form of precipitation and wind impact the formation of lasting freshwater layers in the upper ocean?
- Does precipitation from ARs produce salinity changes in the upper ocean that are detectable by ocean instruments?

This study shows that seasonal freshening in the CCS is linked to cumulative rain fall and AR activity. For locations within 100 km of the coast, the upper 10 m of the CCS freshens throughout the rainy season in response to AR events. Years with higher AR activity are associated with a stronger freshening signal. Studies on event timescales investigate the role that wind plays on the vertical distribution and duration of the freshening signal that results from precipitation. Results show that higher wind speeds induce mixing that brings the freshwater deeper, and decreases the duration of lasting changes in freshwater at the surface. These studies also show that rain events that are characteristic of ARs in the CCS lead to the formation of long-duration (10–50 h) freshwater layers at the surface. We show that rain events that are characteristic of ARs are detectable with observations of the ocean surface salinity: modeled and observational responses to rain events show changes in salinity that are greater than the measurable limit of current CTD instruments.

While this study has shown that freshwater inputs from rain contribute to variability in ocean surface salinity, the use of a one-dimensional model leaves out information about horizontal advection, upwelling, and runoff. Future work could address these shortcomings by applying a similar analysis approach to output from a three-dimensional ocean model. This would

allow for a more comprehensive budget analysis from which conclusion could be drawn about the relative importance of horizontal advection, upwelling, runoff, and external atmospheric forcing. Additionally, future work could apply this type of analysis to properties other than salinity, for example temperature or biogeochemical properties, to further understand the impact of freshwater exchanges on the upper ocean in the CCS.

5.2 Assessing the viability of machine learning (ML) methods for predicting and understanding Arctic sea-ice dynamics.

Numerical models for sea-ice dynamics are highly complex, and must resolve the physics of what amounts to a multi-phase/multi-physics/multi-scale problem. On the other hand, ML models are not subject to constraining to conservation laws, and thus come with improved computational efficiency and cost. Additionally, ML models are driven by data rather than prescribed physical parameterizations that may have inaccuracies. This provides an opportunity to draw information from the data and learn about emergent behaviors that have not yet been recognized.

This work addresses the viability of using ML as a surrogate to parameterize the dynamical component of ice in a numerical model setting. We also analyze the extent to which outputs from XML methods provide consistent and robust results that explain how and why the ML models made their predictions. Models are built to make one-day predictions of sea-ice velocity from various inputs. We use these models to address the following questions:

- Can a ML model in the form of a convolutional neural network (CNN) make skillful predictions of ice motion on one-day time scales?
- How does the performance of the CNN compare with other classical statistical models?
- How does the spatio-temporal variability in the performance of CNN compare with that of the inputs and other properties related to ice motion?

- Are outputs consistent among various explainable machine learning (XML) methods applied to this particular CNN?
- Does output from the layerwise relevance propagation (LRP) XML method show results that are consistent with other more established methods?

We show that a CNN made skillful predictions of sea-ice velocity on one-day time scales. In comparison to other classical statistical models (PS and LR), the CNN has improved performance and the benefits of incorporating both non-linearities between inputs and spatial information when making prediction. Model performances exhibit spatio-temporal variability. In general, improved model performance is linked to:

- increased bathymetric depth and distance from the coast
- larger mean values of u_a , u_i , and c_i
- larger LR coefficients for u_a and c_i ; smaller LR coefficient for u_i

The CNN outperforms the LR in most cases. We have shown that the following are related to increases in the performance of the CNN over the LR:

- larger distance from coast and greater bathymetric depth
- smaller mean u_a and u_i , and larger mean c_i
- larger LR coefficients for u_a , and smaller LR coefficients for c_i

Interestingly, the CNN typically outperforms the LR in regions where wind speed is the dominant LR coefficient, which suggests that the relationship between wind velocity and ice velocity includes non-linearities that are captured by the CNN (and not the LR), leading to an improved performance. This provides motivation to move forward using XML methods that are applied to the CNN rather than analysis of the LR coefficients.

We analyze local and global implementations of various XML methods. Local methods include perturbation and localized LRP. These methods show consistency in the spatial extent to which each input and analysis location impacts predictions of sea-ice velocity. In other words, the radius of influence of the ‘relevance’ metric for each input is similar for the two methods. However, the spatial structure of the areas of relevance is different for the two models: perturbation tends to extract importance from grid points within a particular radius of each analysis point, while the LRP is more non-localised and exhibits high relevance in regions far from the analysis point. These methods have inconsistent results for determining the relative importance of each of the inputs in predicting the output. Perturbation showed that wind speed was the most relevant predictor of ice motion, while LRP showed that sea-ice concentration was the most relevant.

Domain-integrated XML methods also exhibit varying levels of consistency. Here we analyse the LR coefficients in comparison to permutation feature importance (PFI) and a domain-integrated implementation of LRP. PFI and LR are consistent in determining the relative importance of the various inputs, but have inconsistencies in spatial structure. Both methods show the highest relevance for wind speed in predicting sea-ice speed. However, while LR indicates that the relevance of wind is less important in coastal regions, a result consistent with historical studies (Thorndike and Colony, 1982; Serreze et al., 1989; Kwok et al., 2013; Maeda et al., 2020), PFI shows high relevance for wind speed in determining sea-ice motion throughout the Arctic. LRP outputs are inconsistent with both LR and PFI: they show that ice concentration is the most important predictor and the spatial distribution of relevance for each of the inputs does not match historical results.

We find the LRP method to be extremely sensitive to the methods applied during processing the data for input into the CNN, after the model has been trained and it is being prepared for input to LRP, and during manipulations applied to obtain a presentable form of LRP results after the method is applied. In future work we plan to focus on these intricacies and learn more about the causes of large sensitivity in LRP outputs. Moving forward with LRP is desirable because

of the vast amount of information that it provides in comparison to the other XML methods. For example, for LR we can only obtain one parameter (i.e. metric of explainability) for each duration over which the model is applied, while LRP provides a mapped relevance output for each time step and grid location that the CNN is applied.

5.3 Characterizing changes in the relationship between ice drift and wind.

The rapidly declining sea-ice cover in the Arctic has been linked to changes in the dynamical properties of sea ice as the ice enters a regime of free drift where it is more responsive to wind forcing (Rampal et al., 2009; Spreen et al., 2011; Zhang et al., 2012; Kwok et al., 2013; Carmack et al., 2015; Tandon et al., 2017; Docquier et al., 2017). Observed increases in sea-ice drift speeds have been reported, in addition to increases in the ratio between ice and wind speed, a term referred to as the wind factor (Spreen et al., 2011; Maeda et al., 2020). Studies analyzing changes in the wind factor have been carried out using linear regression model of ice as a function of wind speeds. In this study, we build upon historical results by applying a LR model that predicts ice speed as a function of wind, in addition to previous-day ice speed and concentration. We use this model to ask the following questions:

- How is the wind factor changing in time?
- What are the mechanisms contributing to changes in the wind factor?

We confirm historical results: the wind factor is increasing. We find a positive annual trend in the wind factor during months of both minimum and maximum ice extent. We also find positive annual trends in the LR parameter for previous-day ice speed and ice concentration, suggesting the daily predicting of ice speed are becoming more reliant on persistence and the state of the ice. Along with these trends, we see a decrease in ice concentration and increase in drift speed that are consistent with historical results, and the hypothesized ‘free drift mechanism’

whereby as ice concentration decreases, ice enters a regime of free drift where it is more responsive to wind, which leads to increases in ice drift speeds.

When we analyze the spatial variability of changes in the wind factor, a slightly different story is observed, and local yearly negative trends in ice concentration are not always linked to positive trends in the wind factor. We hypothesize several mechanisms that could play a role in the spatial variability of the relationship between the wind factor and ice concentration. These are: the free drift mechanism, free drift with remote effects, geographical features, grid resolution, and the relative importance of ocean currents versus wind speed for driving local ice motion. We discuss implications of each, and propose future methods to evaluate their validity.

In the future we will apply the analysis done for the LR model in chapter 4 to the CNN using LRP as the XML method. In chapter 2 we showed the CNN has improved performance over the LR model. The CNN is able to capture spatial interactions of the inputs, as well as non-linear relationships that exist between the inputs and sea-ice motion. Additionally, XML methods applied to a CNN can provide more information about what is relevant in driving ice motion. This is due to the nature of the output of LRP with CNN compared to LR. For LR we only get one parameter of explainability for each time duration that we run the LR; for example if we build one LR equation for each month/season we get a LR coefficient that explains the relationship between each input and the output for the given month/season. On the other hand, the LRP method applied to a CNN provides a parameter of explainability for each time step that the model is run. Additionally, the LRP provides an entire relevance heat map for each grid point (i.e. for each grid point that a prediction is made the LRP shows a map for what the CNN considered relevant in making a prediction). Therefore, we will apply a CNN with LRP to gain more information about changes to the wind factor (i.e. the relevance of wind in predicting ice motion) than can be provided by the studies using LR. We provide a description of the methods we will use to carry out these studies in section 4.6.2.

For this case we will apply a CNN that has been trained against a baseline reference value, as discussed in chapter 3.4.2. The use of a baseline changes the CNN prediction question

from ‘how important are each of the inputs in predicting ice motion different from zero?’ to ‘how important are each of the inputs in predicting ice motion to be different from the baseline?’ (Mamalakis et al., 2023). This is an important distinction, because here we aim to understand how the relationship between ice motion and wind speed is changing as the ice cover decreases. Therefore, the baseline reference value is chosen to be 0.8% of the wind speed and at an angle of 25.3° to the wind, based on the historical linear relationship between ice speed and wind speed (Thorndike and Colony, 1982). The reference value will be applied to the output for each time step and grid point.

The LRP can also be applied to produce relevance scores with units (Letzgus et al., 2021). We will use the following procedure to train a CNN and run LRP on the outputs. In summary: (i) Subtract the baseline reference value from the outputs; (ii) Standardize the inputs to zero-mean and one standard deviation, but leave the outputs in their original units; (iii) Train the CNN to predict the new output from the inputs; (iv) Un-standardize the inputs (i.e. return to original units); (v) Run LRP at 219 evenly distributed grid points throughout the Arctic. Each of these 219 runs will produce a relevance heat map for each input, indicating the where the input is important for predicting the output at each analyzed location; (vi) Integrate the LRP results over the 219 grid points (i.e. take the average of the 219 maps) to produce a global relevance heat map. We will apply the LRP for the entire duration of the data set, and produces a relevance heat map for each input at each day from 1989–2021.

We also plan to investigate the relationship between the wind factor and ice thickness. We will compare our findings to ice thickness data from CryoSat-2 sea-ice thickness record that has been extended back in time using machine learning techniques (Landy et al., 2022). The mechanism by which ice becomes more responsive to wind forcing is based on decreases in the ice stresses, which are a function of ice concentration, thickness, roughness, and wind speed. Additionally, information about ridges and leads in the ice (i.e. areas of diverging and converging ice) is important for determining internal stresses in the ice (Feltham, 2008). The spatial extent of these linear kinematic features (LKF) in the ice (Hutter et al., 2019) will also be analyzed in

relation to the spatial extent of changes in the wind factor.

5.4 Interdisciplinarity

This research is interdisciplinary in nature. It draws from concepts in physical oceanography and engineering to achieve the overarching goal of gaining a further understanding of the processes involved during exchanges of mass, momentum and energy between the ocean and atmosphere. The first chapter investigates how large precipitation events impact the physical state of the ocean using both oceanic observations and a one-dimensional model with vertical transfer equations for heat and momentum (the MITgcm). The second, third, and fourth chapters involve engineering of a neural network (NN) that, in the future, can be used as a surrogate model for sea-ice motion in physics-based ocean-ice simulations and for skillful forecasting of sea-ice motion. These chapters investigate the physics behind the relationship between sea-ice motion and other atmospheric and oceanic properties, leveraging a NN. Because we are dealing with a multi-phase/multi-physics/multi-scale problem, this work involved an in-depth understanding of the engineering principles of fluid dynamics, mass transport, and heat transport.

Appendix A

Supplemental Material for “Ocean Surface Salinity Response to Atmospheric River Precipitation in the California Current System”

Contents of Appendix A

1. Figures A1 to A5

Introduction This supporting information provides additional figures for analysis of observational data and model outputs (Figs. A1–A5).




Figure A.1. Example MITgcm simulation used for four-day sensitivity tests with idealized forcing from day -2 to day 2. (a) Time series of wind (m s^{-1} , red) and rain (mm hr^{-1} , blue) forcing. (b) Salinity (psu) response in the upper 20 m. The thin black line indicates the depth of the fresh lens, D_L , determined as the depth at which the salinity anomaly relative to the initial time step is 25% of the maximum anomaly for the given run. (c) Salinity difference, ΔS , calculated as the 0.01 m salinity subtracted from the salinity at the first time step. Showing the full length of the four-day simulation, where day zero represents the day with the peak in rain rate and the other date labels are days relative to that day.




Figure A.2. Example of external forcing from ERA5 and initial conditions from Spray for the case study with an AR event starting on October 17, 2016. (a) rain rate (mm hr⁻¹) and wind speed (m s⁻¹), (b) short and longwave radiation (W m⁻²), (c) specific humidity (kg kg⁻¹), and atmospheric temperature (°C), (d) initial salinity (psu) depth profile, and (e) initial temperature (°C) depth profile.




Figure A.3. Comparison of MITgcm vs. GOTM model output of maximum salinity difference, ΔS_{max} , for event-based studies using conditions in the tropics and wind speeds of 4–10 m s⁻¹. The black line represents a linear regression between the two sets of model results, with slope and r^2 labeled in the legend. Outputs for ΔS_{max} using wind speeds below 4 m s⁻¹ are omitted because a 1:1 linear fit is not found in this wind-speed range, as MITgcm produced much higher values for ΔS_{max} than GOTM. GOTM model output was taken from Drushka et al. (2016). It should be noted that here ΔS_{max} is defined as done in Drushka et al. (2016) to allow for a direct comparison between MITgcm and GOTM. Here, ΔS_{max} is the salinity maximum vertical salinity difference between 0.01 m and 5 m for the duration of the event.




Figure A.4. Collection of 85 rain events used in the composite analysis at the MBARI M1 Mooring location from day -3 to day 7. (a) Daily cumulative rain (mm) from ERA5, (b) hourly wind speed (m s^{-1}) with a six hour moving mean from ERA5 and (c) surface salinity anomaly (psu) from day zero as measured by the MBARI M1 Mooring with a one hour moving mean. The different color lines represent the 85 different events that were averaged together for the composite analysis. Day zero is defined to start on the first date with rainfall exceeding a threshold of 5 mm day^{-1} .

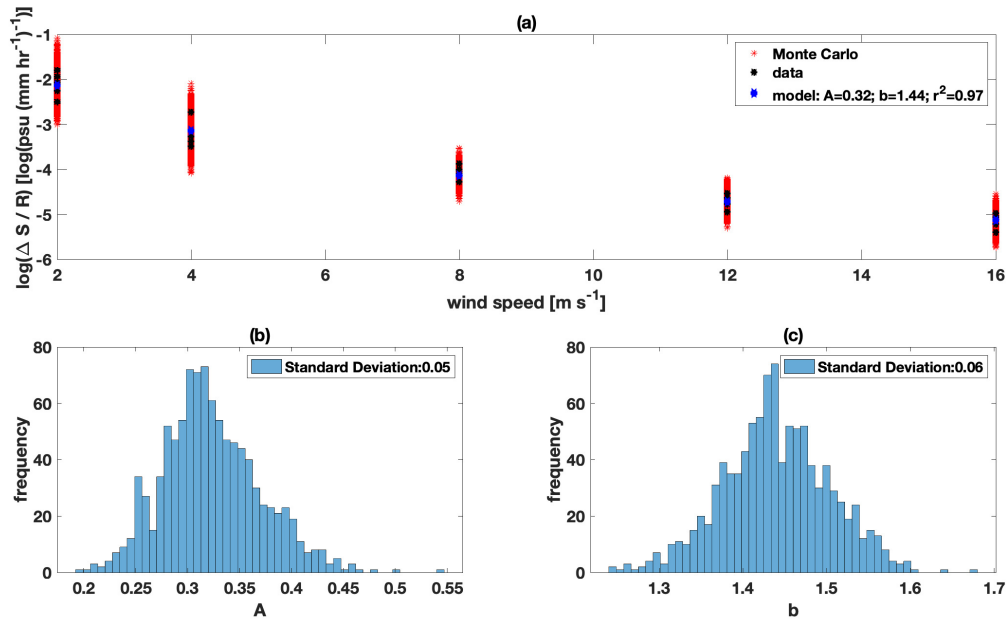


Figure A.5. Monte Carlo simulation to find uncertainty on the parameters A and b in the linear regression equation relating rain rate and precipitation to change in salinity, $\Delta S_{max} = AR_{max}U^b$. In (a) data from the one-dimensional model output (black) for various wind and rain conditions and the linear regression fit of the model output (blue) are compared to the outputs from various linear regression models formed by Monte Carlo simulations (red). The linear regression parameters, A and b , are calculated for each of the 1000 Monte Carlo outputs, and the frequency distributions are shown in (b) and (c).

Appendix B

Supplemental Material for “Machine learning for daily forecasts of Arctic sea-ice motion: an attribution assessment of model predictive skill.”

Contents of Appendix B

1. Figures B1 to B2
2. Tables B1 to B2

Introduction This supporting information provides additional figures for analysis of observational data and model outputs.

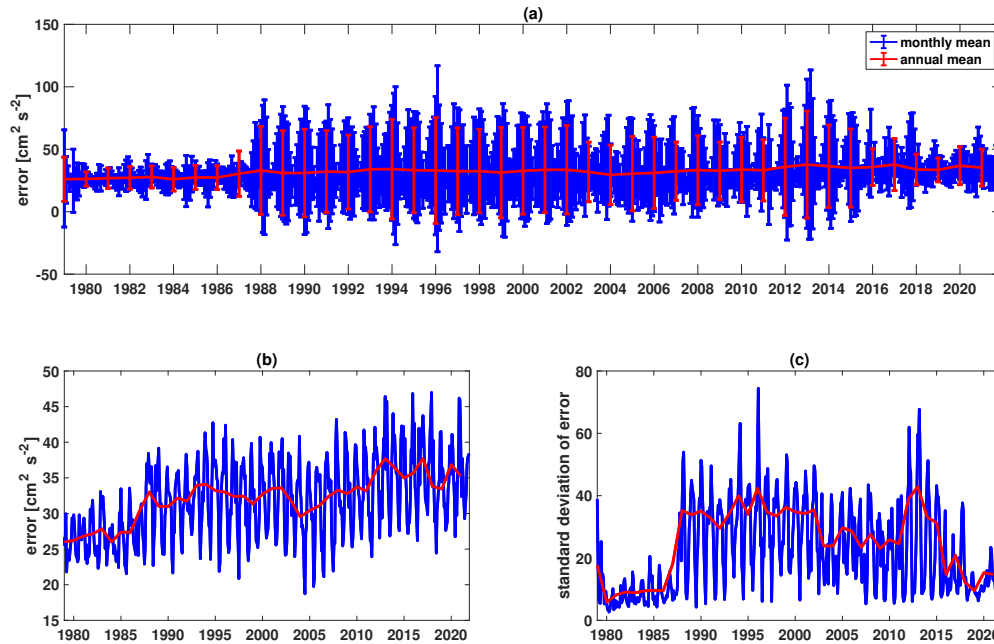


Figure B.1. Uncertainty metrics for the Polar Pathfinder Sea Ice Motion Vector, Version 4 data set. Monthly (blue) and annual (red) mean (a) error of sea-ice velocity with error bars (b) error of sea-ice velocity without error bars, and (c) standard deviation of the error in sea-ice velocity (i.e. error bars in (a)). We note that these uncertainties are not strictly direct error measurements of ice motion. These values come from the relative quality of the source data and the spatial proximity and distribution around each interpolated grid point. A low error means there are (i) higher quality observations, (ii) observations that are closer to the grid cell, (iii) more observations that are within the vicinity of the grid cell that are interpolated, or a combination of all three. The increase in error in Summer 1987 is due to a difference in the sampling period when switching from using Scanning Multichannel Microwave Radiometer (SMMR, 48hr sampling period) to Special Sensor Microwave/Imagers (SSM/Is, 24hr sampling period) for brightness temperature (Tschudi et al. 2020). This inconsistency in the stability of the observing system justifies our use of data from 1989-2022 only.

Table B.1. Comparison of model data and architecture to Zhai and Bitz (2021).

		Hoffman et al (2023)	Zhai and Bitz (2021)
Training Data	Inputs	Present-day wind velocity from JRA55-do, $u_{a,t}$ & $v_{a,t}$ Previous-day sea-ice velocity from Polar Pathfinder Sea Ice Motion, Version 4, $u_{i,t-1}$ & $v_{i,t-1}$ Previous-day sea-ice concentration from Nimbus-7 Passive Microwave, Version 1, c_{t-1}	Present-day wind velocity from JRA55, $u_{a,t}$ & $v_{a,t}$ Previous-day sea-ice velocity from Polar Pathfinder Sea Ice Motion, Version 4, $u_{i,t-1}$ & $v_{i,t-1}$ Previous-day sea-ice concentration from Nimbus-7 Passive Microwave, Version 1, c_{t-1}
	Outputs	Present-day sea-ice velocity from Polar Pathfinder Sea Ice Motion, Version 4, $u_{i,t}$ & $v_{i,t}$	Present-day sea-ice velocity from Polar Pathfinder Sea Ice Motion, Version 4, $u_{i,t}$ & $v_{i,t}$
	Input Size	361 x 361 x 5	40 x 640 x 5
	Output Size	361 x 361 x 2	40 x 640 x 2
Model Architecture	Layers 1–5	Conv2D with ReLU: filter size = (fc,fc), stride = (sc,sc); Max Pool: filter size = (fm,fm), stride = (sm,sm); number of filters = nf	Conv2D with ReLU: filter size = (fc,fc), stride = (sc,sc); Max Pool: filter size = (fm,fm), stride = (sm,sm); number of filters = nf
	Layer 1	fc = (3,3), sc = (1,1); fm = (2,2), sm = (2,2); nf = 7	fc(2,3), sc = (1,1); fm = (2,3), sm = (2,3); nf = 12
	Layer 2	fc = (3,3), sc = (1,1); fm = (2,2), sm = (2,2); nf = 14	fc = (2,3), sc = (1,1); fm = (2,3), sm = (2,3); nf = 24
	Layer 3	fc = (3,3), sc = (1,1); fm = (2,2), sm = (2,2); nf = 28	fc = (2,3), sc = (1,1); fm = (2,2), sm = (2,2); nf = 48
	Layer 4	fc = (3,3), sc = (1,1); fm = (2,2), sm = (2,2); nf = 56	fc = (3,3), sc = (1,1); fm = (2,2), sm = (2,2); nf = 96
	Layer 5	fc = (2,3), sc = (1,1); fm = (2,2), sm = (2,2); nf = 112	fc = (2,3), sc = (1,1); fm = (2,2), sm = (2,2); nf = 192
	Layer 6	Dropout, 20%	Dropout, 20%
	Layer 7	Flatten	Flatten
	Layer 8	Dense	Dense
	Optimizer	Adam	Adam
	Activation	ReLU (slope coefficient, $\alpha = 0.1$)	LeakyReLU ($\alpha = 0.1$)
	Loss	Norm Root Mean Square Error	Norm Root Mean Square Error
	Regularizer	L2 with $\lambda = 0.01$	n/a
	Epochs	50	50
	Batch Size	365	365
Train, Validate, & Test	1989-2017 2018-2019 2020-2021	1990-2014 2015-2016 2017-2018	

Table B.2. Comparison of performance to Zhai and Bitz (2021); calculated from the temporal evaluations (Fig. 5a) for consistency with Fig. 3 in of Zhai and Bitz (2021).

	Hoffman et al (2023)	Zhai and Bitz (2021)
correlation	0.80 ± 0.01	0.82 ± 0.11
skill	0.40 ± 0.02	0.42 ± 0.14

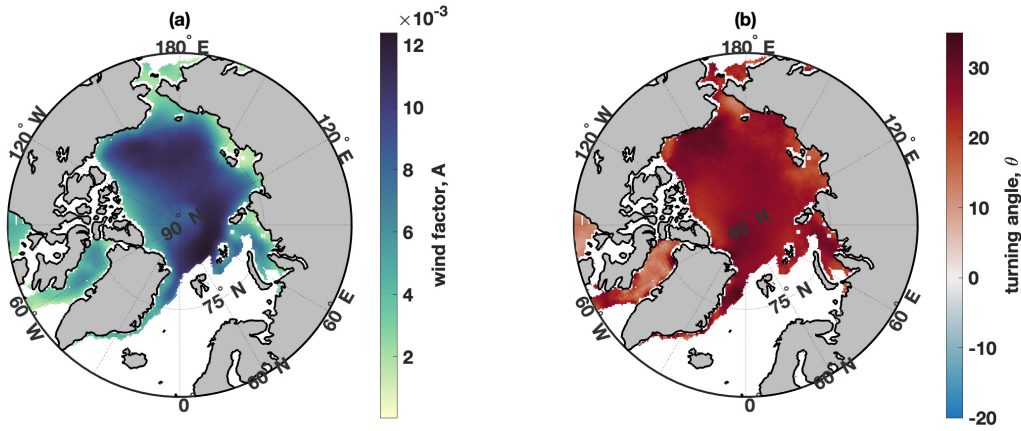


Figure B.2. Ensemble mean (a) wind factor and (b) turning angle for the relationship between wind velocity and sea-ice velocity calculated from the linear regression equation: $u_{i,t}^* = Au_{a,t}^* + B$. Data was not standardized in this case because here we represent the wind factor as the ratio between ice speed and wind speed. The spatially averaged wind factor is $0.72\% \pm 0.31\%$, and the spatially averaged wind angle is $23.3^\circ \pm 6.9^\circ$, which are consistent with historical results from Thorndike and Colony (1982). This wind angle differs slightly from the one cited in the text because it is obtained from a different LR model.

Appendix C

Supplementary Material for “Evaluating the robustness of explainable machine learning (XML) methods for application in regression predictions of Arctic sea-ice motion.”

Contents of Appendix C

1. Figures C1 to C17

Introduction This supporting information provides additional figures for analysis of local XML methods.




Figure C.1. Results from localized sensitivity studies for each of the inputs at location 1, indicated by the red dot. The columns represent each of the different inputs: (a), (d), and (g), wind speed, u_a ; (b), (e), and (h), ice speed, u_i ; and (c), (f), and (i), ice concentration, c_i). The rows represent the different sensitivity methods: (a–c) Spatial covariance; (d–f) normalized RMSE from perturbation analysis; (g–i) normalized relevance score from LRP. Each row is normalized by dividing by the maximum ‘relevance’ value of the spatial mean for each method.




Figure C.2. Same as Fig. C.1, but for location 2.




Figure C.3. Same as Fig. C.1, but for location 3.




Figure C.4. Same as Fig. C.1, but for location 4.




Figure C.5. Same as Fig. C.1, but for location 5.




Figure C.6. Same as Fig. C.1, but for location 6.




Figure C.7. Same as Fig. C.1, but for location 7.




Figure C.8. Same as Fig. C.1, but for location 8.




Figure C.9. Same as Fig. C.1, but for location 9.




Figure C.10. Same as Fig. C.1, but for location 10.




Figure C.11. Same as Fig. C.1, but for location 11.




Figure C.12. Same as Fig. C.1, but for location 12.




Figure C.13. Same as Fig. C.1, but for location 13.

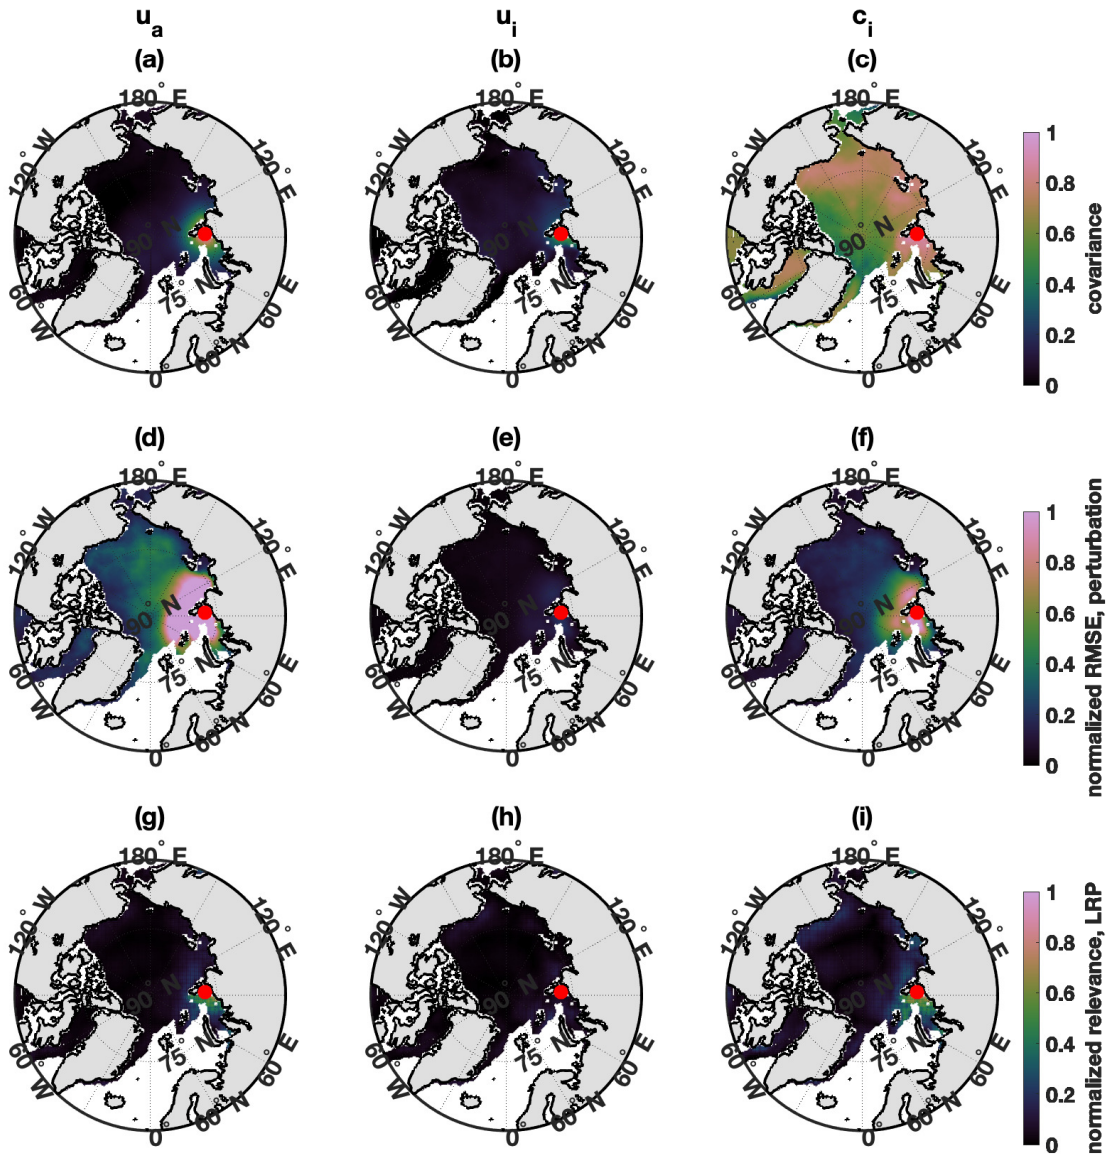


Figure C.14. Same as Fig. C.1, but for location 14.




Figure C.15. Same as Fig. C.1, but for location 15.




Figure C.16. Same as Fig. C.1, but for location 16.




Figure C.17. Same as Fig. C.1, but for location 17.

Bibliography

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. (2015). TensorFlow: Large-scale machine learning on heterogeneous systems. Software available from tensorflow.org, Version 2.8.0.
- Adcroft, A., Campin, J.-M., Dutkiewicz, S., Evangelinos, C., Ferreira, D., Forget, G., Fox-Kepner, B., Heimbach, P., Hill, C., Hill, E., Hill, H., Jahn, O., Losch, M., Marshall, J., Maze, G., Menemenlis, D., and Molod, A. (2018).
- Alber, M., Lapuschkin, S., Seegerer, P., Hagele, M., Schutt, K. T., Montavon, G., Samek, W., Müller, K.-R., Dahne, S., and Kindermans, P.-J. (2019). Investigate neural networks! *Journal of Machine Learning Research*, 20.
- Amador, A., Merrifield, S. T., McCarthy, R. A., Young, R., and Terrill, E. J. (2021). Wave glider speed model for real-time motion planning. In *OCEANS 2021: San Diego – Porto*, pages 1–9.
- Andersson, T. R., Hosking, J. S., Pérez-Ortiz, M., Paige, B., Elliott, A., Russell, C., Law, S., Jones, D. C., Wilkinson, J., Phillips, T., Byrne, J., Tietsche, S., Sarojini, B. B., Blanchard-Wrigglesworth, E., Aksenov, Y., Downie, R., and Shuckburgh, E. (2021). Seasonal Arctic sea ice forecasting with probabilistic deep learning. *Nature Communications*, 12:5124.
- Asher, W. E., Jessup, A. T., Branch, R., and Clark, D. (2014). Observations of rain-induced near-surface salinity anomalies. *Journal of Geophysical Research Oceans*, 119(8):5483–5500.
- Auad, G., Roemmich, D., and Gilson, J. (2011). The California Current System in relation to the Northeast Pacific Ocean circulation. *Progress in Oceanography*, 91(4):576–592.
- Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.-R., and Samek, W. (2015). On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS ONE*, 10(7).

- Bennett, M. M., Stephenson, S. R., Yang, K., Bravo, M. T., and De Jonghe, B. (2020). The opening of the transpolar sea route: Logistical, geopolitical, environmental, and socioeconomic impacts. *Marine Policy*, 73:104178.
- Blanchard-Wrigglesworth, E., Cullather, R., Wang, W., Zhang, J., and Bitz, C. (2015). Model forecast skill and sensitivity to initial conditions in the seasonal sea ice outlook. *Geophysical Research Letters*, 42(19):8042–8048.
- Bograd, S. J., Chereskin, T. K., and Roemmich, D. (2001). Transport of mass, heat, salt, and nutrients in the southern California Current System: Annual cycle and interannual variability. *Journal of Geophysical Research Oceans*, 106(C5):9255–9275.
- Boutin, J., Martin, N., Reverdin, G., Morisset, S., Yin, X., Centurioni, L., and Reul, N. (2014). Sea surface salinity under rain cells: SMOS satellite and in situ drifters observations. *Journal of Geophysical Research Oceans*, 119(8):5533–5545.
- Boutin, J., Martin, N., Reverdin, G., Yin, X., and Gaillard, F. (2013). Sea surface freshening inferred from smos and argo salinity: impact of rain. *Ocean Science*, 9:183–192.
- Brainerd, K. and Gregg, M. (1997). Turbulence and stratification on the Tropical Ocean-Global Atmosphere-Coupled Ocean-Atmosphere Response Experiment microstructure pilot cruise. *Journal of Geophysical Research: Oceans*, 102(C5):10437–10455.
- Brodzik, M. J., Billingsley, B., Haran, T., Raup, B., and Savoie, M. H. (2012). Ease-grid 2.0: Incremental but significant improvements for earth-gridded data sets. *ISPRS International Journal of Geo-Information*, 1(1):32–45.
- Cao, Y., Liang, S., Sun, L., Liu, J., Cheng, X., Wang, D., Chen, Y., Meng, Y., and Kuishuang, F. (2022). Trans-arctic shipping routes expanding faster than the model predictions. *Global Environmental Change*, 73:102488.
- Carmack, E., Polyakov, I., Padman, L., Fer, I., Hunke, E., Hutchings, J., Jackson, J., Kelley, D., Kwok, R., Layton, C., Melling, H., Perovich, D., Persson, O., Ruddick, B., Timmermans, M.-L., Toole, J., Vavrus, S., and Winsor, P. (2015). Toward quantifying the increasing role of oceanic heat in sea ice loss in the new Arctic. *Bulletin of the American Meteorological Society*, 96(12):2079–2105.
- Cavalieri, D., Parkinson, C., Gloersen, P., and Zwally, H. (1996). Sea ice concentrations from Nimbus-7 SMMR and DMSP SSM/I-SSMIS Passive Microwave Data, Version 1. Access Data: March 2021.
- Chaudhuri, D., Sengupta, D., D’Asaro, E., and S., S. (2021). Trapping of wind momentum in a salinity-stratified ocean. *Journal of Geophysical Research: Oceans*, 126(e2021JC017770).

- Chavez, F. P. (2015). In situ, meteorological, physical, and profile data collected by Monterey Bay Aquarium Research Institute at OceanSITES site MBARI from 2004-04-30 to 2021-09-07 (NCEI Accession 0130040). Access Data: December 2020.
- Clayson, C. A., Edson, J. B., Paget, A., Graham, R., and Greenwood, B. (2019). Effects of rainfall on the atmosphere and the ocean during SPURS-2. *Oceanography*, 32(2):86–97.
- Cole, S. T., Toole, J. M., Lele, R., Timmermans, M.-L., Gallaher, S. G., Stanton, T. P., Shaw, W. J., Hwang, B., Maksym, T., Wilkinson, J. P., Ortiz, M., Graber, H., Rainville, L., Petty, A. A., Farrell, S. L., Richter-Menge, J. A., and Haas, C. (2017). Ice and ocean velocity in the Arctic marginal ice zone: Ice roughness and momentum transfer. *Elementa: Science of the Anthropocene*, 5:55.
- Crawford, A., Stroeve, J., Smith, A., and Jahn, A. (2021). Arctic open-water periods are projected to lengthen dramatically by 2100. *Communications Earth & Environment*, 2(1):109.
- Davis, R. E., Ohman, M. D., Rudkick, D. L., Sherman, J. T., and Hodges, B. (2008). Glider surveillance of physics and biology in the southern California Current System. *Limnol. Oceanography*, 53:2151–2168.
- de Boyer Montégut, C., Mignot, J., Lazar, A., and Cravatte, S. (2007). Control of salinity on the mixed layer depth in the world ocean: 1. General description. *Journal of Geophysical Research*, 112(C06011).
- Delcroix, T., McPhaden, M. J., Dessier, A., and Gouriou, Y. (2005). Time and space scales for sea surface salinity in the tropical oceans. *Deep Sea Research Part I: Oceanographic Research Papers*, 52(5):787–813.
- Dettinger, M. (2011). Climate change, atmospheric rivers, and floods in California — a multi-model analysis of storm frequency and magnitude changes. *J. Amer. Water Resour. Assoc.*, 47.
- Dietterich, T. (1998). Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Computation*, 10:1895–1923.
- Dobrescu, A., Giuffrida, M. V., and Tsafaris, S. A. (2019). Understanding deep neural networks for regression in leaf counting. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2600–2608.
- Docquier, D., Massonnet, F., Barthelemy, A., Tandon, N. F., Lecomte, O., and Fichet, T. (2017). Relationships between arctic sea ice drift and strength modelled by nemo-lim3.6. *Cryosphere*, 11:2829–2846.
- Drucker, R. and Riser, S. C. (2014). Validation of Aquarius sea surface salinity with Argo:

- Analysis of error due to depth of measurement and vertical salinity stratification. *Journal of Geophysical Research Oceans*, 119(7):4626–4637.
- Drushka, K., Asher, W. E., Jessup, A. T., Thompson, E. J., Iyer, S., and Clark, D. (2019). Capturing fresh layers with the surface salinity profiler. *Oceanography*, 32(2):76–85.
- Drushka, K., Asher, W. E., Ward, B., and Walesby, K. (2016). Understanding the formation and evolution of rain-formed fresh lenses at the ocean surface. *JGR Oceans*, 121(4):2673–2689.
- Dumitru, C. O., Andrei, V., Schwarz, G., and Datcu, M. (2019). Machine learning for sea ice monitoring from satellites. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLII-2/W16:83–89.
- Ekman, W. V. (1905). On the influence of the earth's rotation on ocean-currents. *Arkiv for Matematik, Astronomi Och Fysik*, 2(11):1–52.
- Feltham, D. L. (2008). Sea ice rheology. *Annual Review of Fluid Mechanics*, 40(1):91–112.
- Fish, M. A., Wilson, A. M., and Ralph, M. F. (2019). Atmospheric river families: Definition and associated synoptic conditions. *Journal of Hydrometeorology*, 21(3).
- Gershunov, A. (2017). Catalog of landfalling atmospheric rivers along the western coast of North America. Access Data: December 2020.
- Gershunov, A., Shulgina, T., Ralph, F. M., Lavers, D. A., and Rutz, J. J. (2017). Accessing the climate-scale variability of atmospheric rivers affecting western North America. *Geophysical Research Letters*, 44(15):7900–7908.
- Giglio, D., Lyubchich, V., and Gille, S. T. (2020). Seasonal to interannual variability of upper ocean temperature and salinity: The role of atmospheric rivers. Number AI14A-2255, San Diego. Ocean Sciences Meeting.
- Goosse, H. and Fichefet, T. (1999). Importance of ice-ocean interactions for the global ocean circulation: A model study. *Journal of Geophysical Research: Oceans*, 104(C10):23337–23355.
- Guan, B. and Waliser, D. E. (2015). Detection of atmospheric rivers: Evaluation and application of an algorithm for global studies. *J. Geophys. Res. Atmos.*, 120:12514 – 12535.
- Haar, L. V., Elvira, T., and Ochoa, O. (2023). An analysis of explainability methods for convolutional neural networks. *Engineering Applications of Artificial Intelligence*, 117:105606.
- Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz-Sabater, J., Nicolas, J., Pubeby, C., Radu, R., Schepers, D., Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo,

- G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., De Chiara, G., Dahlgren, P., Dee, D., Diamantakis, M., Dragon, R., Flemming, J., Forbes, R., Fuentes, M., Geer, A., Haimberger, L., Healy, S., Hogan, R. J., Hólm, E., Janisková, M., Keeley, S., Laloyaux, P., Lopez, P., Lupu, C., Radnoti, G., de Rosnay, P., Rozum, I., Vamborg, F., Villaume, S., and Thépaut, J.-N. (2020). The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146(730):1999–2049.
- Hibler, W. (1979). A dynamic thermodynamic sea ice model. *Journal of Physical Oceanography*, 9(4):815–846.
- Hoffman, L., Mazloff, M. R., Gille, S. T., Giglio, D., Bitz, C. M., Heimbach, P., and Matsuyoshi, K. (2023). Data from: Machine learning for daily forecasts of arctic sea-ice motion: an attribution assessment of model predictive skill. Publish Data: June 2023.
- Hsieh, W. W. and Tang, B. (1998). Applying neural network models to prediction and data analysis in meteorology and oceanography. *Bulletin of the American Meteorological Society*, 79(9):1855 – 1870.
- Hunke, E., Allard, R., Blain, P., Blockley, E., Feltham, D., Fichefet, T., Garric, G., Grumbine, R., Lemieux, J.-F., Rasmussen, T., Ribergaard, M., Roberts, A., Schweiger, A., Tietsche, S., Tremblay, B., Vancoppenolle, M., and Zhang, J. (2020). Should sea-ice modeling tools designed for climate research be used for short-term forecasting? *Current Climate Change Reports*, 6:121–136.
- Hunter, J. D. (2007). Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95.
- Hutter, N., Zampieri, L., and Losch, M. (2019). Linear kinematic features (leads & pressure ridges) detected and tracked in radarsat geophysical processor system (rgps) sea-ice deformation data from 1997 to 2008.
- Ivanovs, M., Kadikis, R., and Ozols, K. (2021). Perturbation-based methods for explaining deep neural networks: A survey. *Pattern Recognition Letters*, 150:228–234.
- Iyer, S. and Drushka, K. (2021). The influence of preexisting stratification and tropical rain modes on the mixed layer salinity response to rainfall. *Journal of Geophysical Research Oceans*, 126(e2021JC017574).
- Jahn, A. (2018). Reduced probability of ice-free summers for 1.5°C compared to 2°C warming. *Nature Climate Change*, 8:409–413.
- Jakobsson, M., Mayer, L. A., Bringensparr, C., Castro, C. F., Mohammad, R., Johnson, P., Ketter, T., Accettella, D., Amblas, D., An, L., Arndt, J. E., Canals, M., Casamor, J. L., Chauché, N., Coakley, B., Danielson, S., Demarte, M., Dickson, M.-L., Dorschel, B., Dowdeswell,

- J. A., Dreutter, S., Fremand, A. C., Gallant, Dana and Hall, J. K., Hehemann, L., Hodnesdal, H., Hong, J., Ivaldi, R., Kane, Emily and Klaucke, I., Krawczyk, D. W., Kristoffersen, Yngve and Kuipers, B. R., Millan, R., Masetti, Giuseppe and Morlighem, M., Noormets, R., Prescott, M. M., Rebesco, M., Rignot, E., Semiletov, I., Tate, A. J., Travaglini, P., Velicogna, I., Weatherall, P., Weinrebe, W., Willis, J. K., Wood, M., Zarayskaya, Yulia and Zhang, T., Zimmermann, M., and Zinglensen, K. B. (2020). The international bathymetric chart of the arctic ocean version 4.0. *Sci Data*, 7(176).
- Jeffries, M. O., Overland, J. E., and Perovich, D. K. (2013). The Arctic shifts to a new normal. *Physics Today*, 66(10):35–40.
- Johnson, K., Chavez, F., and Friederich, G. (1999). Continental-shelf sediment as a primary source of iron for coastal phytoplankton. *Nature*, 398:697–700.
- Kern, S., Lavergne, T., Notz, D., Pedersen, L. T., Tonboe, R. T., Saldo, R., and Sørensen, A. M. (2019). Satellite passive microwave sea-ice concentration data set intercomparison: closed ice and ship-based observations. *The Cryosphere*, 13(12):3261–3307.
- Kim, Y. J., Kim, H.-C., Han, D., Lee, S., and Im, J. (2020). Prediction of monthly arctic sea ice concentrations using satellite and reanalysis data based on convolutional neural networks. *The Cryosphere*, 14(3):1083–1104.
- Kimura, N., Nishimura, A., Tanaka, Y., and Yamaguchi, H. (2013). Influence of winter sea-ice motion on summer ice cover in the arctic. *Polar Research*, 33(20193).
- Kimura, N. and Wakatsuchi, M. (2000). Relationship between sea-ice motion and geostrophic wind in the northern hemisphere. *Geophysical Research Letters*, 27(22):3735–3738.
- Kobayashi, S., Ota, Y., Harada, Y., Ebata, A., Moriya, M., Onoda, H., Onogi, K., Kamahori, H., Kobayashi, C., Endo, H., Miyaoka, K., and Takahashi, K. (2015). The jra-55 reanalysis: General specifications and basic characteristics. *Journal of the Meteorological Society of Japan. Ser. II*, 93(1):5–48.
- Kudela, R. and Chavez, F. (2004). The impact of coastal runoff on ocean color during an El Niño year in Central California. *Deep-Sea Research II*, 51(10-11):1173–1185.
- Kwok, R. (2018). Arctic sea ice thickness, volume, and multiyear ice coverage: losses and coupled variability (1958–2018). *Environmental Research Letters*, 13(10).
- Kwok, R., Spreen, G., and Pang, S. (2013). Arctic sea ice circulation and drift speed: Decadal trends and ocean currents. *Journal of Geophysical Research: Oceans*, 118(5):2408–2425.
- Lagerloef, G., SCHMITT, R., SCHANZE, J., and KAO, H.-Y. (2010). The ocean and the global water cycle. *Oceanography*, 23(4):82–93.

- Lakshmanan, V., Karstens, C., Krause, J., Elmore, K., Ryzhkov, A., and Berkseth, S. (2015). Which polarimetric variables are important for weather/no-weather discrimination? *Journal of Atmospheric and Oceanic Technology*, 32(6):1209 – 1223.
- Landy, J., Dawson, G., and Tsamados, M. (2022). A year-round satellite sea-ice thickness record from cryosat-2. *Nature*, 609:517–522.
- Large, W., McWilliams, J., and Doney, S. (1994). Ocean vertical mixing: A review and a model with a nonlocal boundary layer parameterization. *Reviews of Geophysics*, 32(4):363–403.
- Large, W. and Pond, S. (1982). Sensible and latent heat flux measurements over the ocean. *Journal of Physical Oceanography*, 12(5):464–482.
- Lee, S., Im, J., Kim, J., Kim, M., Shin, M., Kim, H.-c., and Quackenbush, L. J. (2016). Arctic sea ice thickness estimation from cryosat-2 satellite data using machine learning-based lead detection. *Remote Sensing*, 8(9).
- Lepparanta, M. (2011). *The Drift of Sea Ice*. Springer-Praxis Books in Geophysical Science. Springer-Verlag, Berlin Heidelberg.
- Letzgus, S., Wagner, P., Lederer, J., Samek, W., Müller, K., and Montavon, G. (2021). Toward explainable AI for regression models. *CoRR*, abs/2112.11407.
- Li, M., Zhang, R., and Liu, K. (2021). Machine learning incorporated with causal analysis for short-term prediction of sea-ice. *Frontiers in Marine Science*, 8(649378).
- Linardatos, P., Papastefanopoulos, V., and Kotsiantis, S. (2021). Explainable ai: A review of machine learning interpretability methods. *Entropy*, 23(1).
- Lynn, R. J. and Simpson, J. J. (1987). The California Current System: The seasonal variability of its physical characteristics. *Journal of Geophysical Research*, 92(C12):12947–12966.
- Maeda, K., Kimura, N., and Yamaguchi, H. (2020). Temporal and spatial change in the relationship between sea-ice motion and wind in the arctic. *Polar Research*, 39.
- Mamalakis, A., Barnes, E. A., and Ebert-Uphoff, I. (2023). Carefully choose the baseline: Lessons learned from applying xai attribution methods for regression tasks in geoscience. *Artificial Intelligence for the Earth Systems*, 2(1):e220058.
- Marquardt, D. W. and Snee, R. D. (1975). Ridge regression in practice. *The American Statistician*, 29(1):3–20.
- Maykut, G., Thorndike, A., and Untersteiner, N. (1972). Aidx scientific plan. *AIDJEX Bulletin*, 15:1–67.

- McCulloch, M., Spurgeon, P., and Chuprin, A. (2012). Have mid-latitude ocean rain-lenses been seen by the SMOS satellite? *Ocean Modelling*, 43-44:108–111.
- McGovern, A., Lagerquist, R., Gagne, D. J., Jergensen, G. E., Elmore, K. L., Homeyer, C. R., and Smith, T. (2019). Making the black box more transparent: Understanding the physical implications of machine learning. *Bulletin of the American Meteorological Society*, 100(11):2175 – 2199.
- Meredith, M., Sommerkorn, M., Cassotta, S. and Derksen, C., Ekaykin, A., Hollowed, A., Kofinas, G., Mackintosh, A., Melbourne-Thomas, J., Muelbert, M., Ottersen, G., Pritchard, H., and Schuur, E. (2019). Chapter 3: Polar regions. in: IPCC special report on the ocean and cryosphere in a changing climate.
- Molina, M. J., Gagne, D. J., and Prein, A. F. (2021). A benchmark to test generalization capabilities of deep learning methods to classify severe convective storms in a changing climate. *Earth and Space Science*, 8(9):e2020EA001490.
- Montavon, G., Bach, S., Binder, A., Samek, W., and Müller, K. (2015). Explaining nonlinear classification decisions with deep taylor decomposition. *CoRR*, abs/1512.02479.
- Montavon, G., Binder, A., Lapuschkin, S., Samek, W., and Muller, K.-R. (2019). Layer-wise relevance propagation: an overview. In Samek, W., Montavon, G., Vedaldi, A., Hansen, L., and Muller, K.-R., editors, *Explainable AI*, volume 11700 of *LNCS*, pages 193–209. Springer Cham.
- Montavon, G., Samek, W., and Muller, K.-R. (2018). Methods for interpreting and understanding deep neural networks. *Digital Signal Processing*, 73:1–15.
- Muñoz Sabater, J. (2019). ERA5-Land hourly data from 1981 to present.
- Notz, D. and Stroeve, J. (2018). The trajectory towards a seasonally ice-free arctic ocean. *Current Climate Change Reports*, 4:407–416.
- Olason, E. and Notz, D. (2014). Drivers of variability in Arctic sea-ice drift speed. *Journal of Geophysical Research: Oceans*, 119(9):5755–5775.
- Onarheim, I. H., Eldevik, T., Smedsrud, L. H., and Stroeve, J. (2018). Seasonal and regional manifestation of arctic sea ice loss. *Journal of Climate*, 13(12):4917–4932.
- O’Shea, K. and Nash, R. (2015). An introduction to convolutional neural networks.
- Payne, A. E., Demory, M.-E., Leung, L. R., Ramos, A. M., Shields, C. A., Rutz, J. J., Siler, N., Villarini, G., Hall, A., and Ralph, F. M. (2020). Responses of atmospheric rivers to climate change. *Nature Reviews Earth and Environment*, 1:143–157.

- Petrou, Z. I. and Tian, Y. (2019). Prediction of sea ice motion with convolutional long short-term memory networks. *IEEE Transactions on Geoscience and Remote Sensing*, 57(9).
- Polyakov, I. V., Rippeth, T. P., Fer, I., Alkire, M. B., Baumann, T. M., Carmack, E. C., Ingvaldsen, R., Ivanov, V. V., Janout, M., Lind, S., Padman, L., Pnyushkov, A. V., and Rember, R. (2020). Weakening of cold halocline layer exposes sea ice to oceanic heat in the eastern arctic ocean. *Journal of Climate*, 33(18):8107 – 8123.
- Press, W., Flannery, B., Teukolsky, S., and Vetterling, W. (1986). *Numerical Recipes*. Cambridge University Press.
- Price, J. F. (1979). Observations of a rain-formed mixed layer. *Journal of Physical Oceanography*, 9(3).
- Radivojac, P., Obradovic, Z., Dunker, A. K., and Vucetic, S. (2004). Feature selection filters based on the permutation test. In Boulicaut, J.-F., Esposito, F., Giannotti, F., and Pedreschi, D., editors, *Machine Learning: ECML 2004*, pages 334–346, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Rahman, M. M., Matsuo, K., Matsuzaki, S., and Purushotham, S. (2021). Deeppseudo: Pseudo value based deep learning models for competing risk analysis. *AAAI Conference on Artificial Intelligence*, 35(1).
- Ralph, F., Coleman, T., Neiman, J., Zamora, J., and Dettinger, D. (2013). Observed impacts of duration and seasonality of atmospheric-river landfalls on soil moisture and runoff in coastal northern California. *Journal of Hydrometeorology*, 14:443–459.
- Ralph, F. and Dettinger, M. (2011). Storms, floods and the science of atmospheric rivers. *Eos Trans. AGU*, 92(32):265–272.
- Ralph, F. and Dettinger, M. (2012). Historical and national perspectives on extreme West Coast precipitation associated with atmospheric rivers during december 2010. *Bull. Amer. Meteor. Soc.*, 93(6):783–790.
- Ramon, J., Lledó, L., Torralba, V., Soret, A., and Doblas-Reyes, F. J. (2019). What global reanalysis best represents near-surface winds? *Quarterly Journal of the Royal Meteorological Society*, 145(724):3236–3251.
- Ramos, A. M., Roca, R., Soares, P. M., Wilson, A. M., Trigo, R. M., and Ralph, F. M. (2021). Uncertainty in different precipitation products in the case of two atmospheric river events. *Environmental Research Letters*, 16(045012).
- Rampal, P., Weiss, J., and Marsan, D. (2009). Positive trend in the mean speed and deformation rate of arctic sea ice, 1979-2007. *Journal of Geophysical Research*, 114(C5):C05013.

- Ren, A. S. and Rudnick, D. L. (2021). Temperature and salinity extremes from 2014-2019 in the California Current System and its source waters. *Nature Communications Earth & Environment*, 2(62).
- Ren, L. and Riser, S. C. (2009). Seasonal salt budget in the northeast Pacific Ocean. *Journal of Geophysical Research*, 114(C12).
- Rudnick, D. (2016). California underwater glider network. Access Data: May 2020.
- Rudnick, D., Zaba, K., Todd, R., and Davis, R. (2017a). A climatology using data from the California Underwater Glider Network. Access Data: May 2020.
- Rudnick, D., Zaba, K. D., Todd, R. E., and Davis, R. E. (2017b). A climatology of the California Current System from a network of underwater gliders. *Progress in Oceanography*, 154:64–106.
- Saludo, R. and Hackett, B. (2020). Global ocean - high resolution sar sea ice drift.
- Schmitt, R. (2008). Salinity and the global water cycle. *Oceanography*, 21(1):12–19.
- Schmitt, R. W. (1995). The ocean component of the global water cycle. *Reviews of Geophysics*, 33(S2):1395–1409.
- Schnake, T., Eberle, O., Lederer, J., Nakajima, S., Schütt, K. T., Müller, K., and Montavon, G. (2020). XAI for graphs: Explaining graph neural network predictions by identifying relevant walks. *CoRR*, abs/2006.03589.
- Schneider, N., Di Lorenzo, E., and Niiler, P. P. (2005). Salinity variations in the southern California Current. *Journal of Physical Oceanography*, 35(8):1421–1436.
- Seabird Scientific (2016). How accurate is salinity measured by my ctd? what factors impact accuracy?
- Serreze, M. C., Barry, R. G., and McLaren, A. S. (1989). Seasonal variations in sea ice motion and effects on sea ice concentration in the Canada basin. *Journal of Geophysical Research: Oceans*, 94(C8):10955–10970.
- Serreze, M. C., Holland, M. M., and Stroeve, J. (2007). Perspectives on the Arctic's shrinking ice cover. *Science*, 315(5815):1533–1536.
- Shen, X., Zhang, J., Zhang, X., Meng, J., and Ke, C. (2017). Sea ice classification using Cryosat-2 altimeter data by optimal classifier–feature assembly. *IEEE Geoscience and Remote Sensing Letters*, 14(11):1948–1952.
- Shields, C. A. and Kiehl, J. T. (2016). Simulating the Pineapple Express in the half degree Com-

- munity Climate Change System Model, CCSM4. *Geophysical Research Letters*, 43(14):7767–7773.
- Sinha, A. and Abernathey, R. (2021). Estimating ocean surface currents with machine learning. *Frontiers in Marine Science*, 8.
- Smyth, W., Zavialov, P., and Moum, J. (1997). Decay of turbulence in the upper ocean following sudden isolation from surface forcing. *Journal of Physical Oceanography*, 27(5):810–822.
- Solomon, A., Heuzé, C., Rabe, B., Bacon, S., Bertino, L., Heimbach, P., Inoue, J., Iovino, D., Mottram, R., Zhang, X., Aksenov, Y., McAdam, R., Nguyen, A., Raj, R. P., and Tang, H. (2021). Freshwater in the arctic ocean 2010–2019. *Ocean Science*, 17(4):1081–1102.
- Soloviev, A. V., Matt, S., and Fujimura, A. (2015). Three-dimensional dynamics of freshwater lenses in the ocean’s near-surface layer. *Oceanography*, 28(1):142–149.
- Sonneveld, M., Lguensat, R., Jones, D. C., Deuben, P. D., Brajard, J., and Balaji, V. (2021). Bridging observations, theory and numerical simulation of the ocean using machine learning. *Environmental Research Letters*, 16(073008).
- Spreen, G., Kwok, R., and Menemenlis, D. (2011). Trends in arctic sea ice drift and role of wind forcing: 1992–2009. *Geophysical Research Letters*, 38(L19501).
- SPURS-2 Planning Group (2015). From salty to fresh—salinity processes in the upper-ocean regional study-2 (SPURS-2): Diagnosing the physics of a rainfall-dominated salinity minimum. *Oceanography*, 28(1):150–159.
- Stephen23 (2023). Cubehelix colormap generator: Beautiful and versatile! Access Date: January 2023.
- Stroeve, J. and Notz, D. (2018). Changing state of arctic sea ice across all seasons. *Environmental Research Letters*, 13(10):103001.
- Stroeve, J., Serreze, M. C., Holland, M. M., Kay, J. E., Malanik, J., and Barrett, A. P. (2012). The Arctic’s rapidly shrinking sea ice cover: A research synthesis. *Climate Change*, 110:1005–1027.
- Sumata, H., Kauker, F., Karcher, M., Rabe, B., Timmermans, M.-L., Behrendt, A., Gerdes, R., Schauer, U., Shimada, K., Cho, K.-H., and Kikuchi, T. (2018). Decorrelation scales for arctic ocean hydrography – part i: Amerasian basin. *Ocean Science*, 14(1):161–185.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. (2015). Going deeper with convolutions. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–9.

- Sévellec, F., Fedorov, A. V., and Liu, W. (2017). Arctic sea-ice decline weakens the atlantic meridional overturning circulation. *Nature Climate Change*, 7:604–610.
- Taboda, F. F., Stock, C. A., Griffies, S. M., Dunne, J., John, J. G., Small, J. R., and Tsujino, H. (2019). Surface winds from atmospheric reanalysis lead to contrasting oceanic forcing and coastal upwelling patterns. *Ocean Modelling*, 133:79–111.
- Tandon, N. F., Kushner, P. J., Docquier, D., Wettstein, J. J., and Li, C. (2017). Reassessing sea ice drift and its relationship to long-term arctic sea ice loss in coupled climate models. *Journal of Geophysical Research: Oceans*, 123:4338–4359.
- Tang, B., Hsieh, W. W., Monahan, A. H., and Tangang, F. T. (2000). Skill comparisons between neural networks and canonical correlation analysis in predicting the equatorial pacific sea surface temperatures. *Journal of Climate*, 13(1):287–293.
- Tarek, M., Brissette, F., and Arsenault, R. (2020). Evaluation of the ERA5 reanalysis as a potential reference dataset for hydrological modelling over North America. *Hydrological Earth Systems Science*, 24:2527–2544.
- Thoman, R. L., Druckenmiller, M. L., Moon, T. A., Andreassen, L. M., Baker, E., Ballinger, T. J., Berner, L. T., Bernhard, G. H., Bhatt, U. S., Bjerke, J. W., Boisvert, L., Box, J. E., Brettschneider, B., Burgess, D., Butler, A. H., Cappelen, J., Christiansen, H. H., Decharme, B., Derksen, C., Divine, D., Drozdov, D. S., Elias, C. A., Epstein, H. E., Farrell, S. L., Fausto, R. S., Fettweis, X., Fioletov, V. E., Forbes, B. C., Frost, G. V., Gerland, S., Goetz, S. J., Grooß, J.-U., Haas, C., Hanna, E., Hanssen, B. I., Heijmans, M. M. P. D., Hendricks, S., Ialongo, I., Isaksen, K., Jensen, C. D., Johnsen, B., Kaleschke, L., Kholodov, A. L., Kim, S.-J., Kohler, J., Korsgaard, N. J., Labe, Z., Lakkala, K., Lara, M. J., Lee, S. H., Loomis, B., Luks, B., Luoju, K., Macander, M. J., Magnússon, R., Malkova, G. V., Mankoff, K. D., Manney, G. L., Meier, W. N., Mote, T., Mudryk, L., Müller, R., Nyland, K. E., Overland, J. E., Pálsson, F., Park, T., Parker, C. L., Perovich, D., Petty, A., Phoenix, G. K., Pinzon, J. E., Ricker, R., Romanovsky, V. E., Serbin, S. P., Sheffield, G., Shiklomanov, N. I., Smith, S. L., Stafford, K. M., Steer, A., Streletskiy, D. A., Svendby, T., Tedesco, M., Thomson, L., Thorsteinsson, T., Tian-Kunze, X., Timmermans, M.-L., Tømmervik, H., Tschudi, M., Tucker, C. J., Walker, D. A., Walsh, J. E., Wang, M., Webster, M., Wehrlé, A., Øyvind Winton, Wolken, G., Wood, K., Wouters, B., and Yang, D. (2022). The arctic. *Bulletin of the American Meteorological Society*, 103(8):S257 – S306.
- Thompson, E. J., Moum, J. N., Fairall, C. W., and Rutledge, S. A. (2019). Wind limits on rain layers and diurnal warm layers. *Journal of Geophysical Research: Oceans*, 124:897–924.
- Thomson, R. E. and Emery, W. J. (2014). *Statistical Methods and Error Handling. Data Analysis Methods in Physical Oceanography* (Third Edition). Elsevier.
- Thorndike, A. and Colony, R. (1982). Sea ice motion in response to geostrophic winds. *Journal*

- of Geophysical Research: Oceans*, 87(C8):5845–5852.
- Thyng, K. M., Greene, C. A., Hetland, R. D., Zimmerle, H. M., and DiMarco, S. F. (2016). True colors of oceanography: Guidelines for effective and accurate colormap selection. *Oceanography*, 29(3).
- Timmermans, M.-L. and Marshall, J. (2020). Understanding arctic ocean circulation: A review of ocean dynamics in a changing climate. *Journal of Geophysical Research: Oceans*, 125(4):e2018JC014378.
- Tomczak, M. (1995). Salinity variability in the surface layer of the tropical western Pacific Ocean. *Journal of Geophysical Research Oceans*, 100(C10):20499–20515.
- Toms, B. A., Barnes, E. A., and Ebert-Uphott, I. (2020). Physically interpretable neural networks for the geosciences: Applications to earth system variability. *Journal of Advances in Modeling Earth Systems*, 12(9):e2019MS002002.
- Tschudi, M., Meier, W., Stewart, J., Fowler, C., and Maslanik, J. (2019). Polar Pathfinder Daily 25 km EASE-Grid Sea Ice Motion Vectors, Version 4. Access Data: March 2021.
- Tschudi, M. A., Meier, W. N., and Steward, J. S. (2020). An enhancement to sea ice motion and age products at the National Snow and Ice Data Center (NSIDC). *The Cryosphere*, 14(5):1519–1536.
- Tsujino, H., Urakawa, S., Nakano, H., Small, R. J., Kim, W. M., Yeager, S. G., Danabasoglu, G., Suzuki, T., Bamber, J. L., Bentsen, M., Böning, C. W., Bozec, A., Chassignet, E. P., Curchitser, E., Dias, F. B., Durack, P. J., Griffies, S. M., Harada, Y., Ilicak, M., Josey, S. A., Kobayashi, C., Kobayashi, S., Komuro, Y., Large, W. G., Le Sommer, J., Marsland, S. J., Masina, S., Scheinert, M., Tomita, H., Valdivieso, M., and Yamazaki, D. (2018). JRA-55 based surface dataset for driving ocean–sea-ice models (JRA55-do). *Ocean Modelling*, 130:79–139.
- Untersteiner, N., Thorndike, A., Rothrock, D., and Hunkins, K. (2007). AIDJEX Revisited: A look back at the U.S.-Canadian Arctic Ice Dynamics Joint Experiment 1970-78. *Arctic*, 60(3):327–336.
- Vinogradova, N., Lee, T., Boutin, J., Drushka, K., Fournier, S., Sabia, R., Stammer, D., Bayler, E., Reul, N., Gordon, A., Melnichenko, O., Li, L., Hackert, E., Martin, M., Kolodziejczyk, N., Hasson, A., Brown, S., Misra, S., and Lindstrom, E. (2019). Satellite salinity observing system: Recent discoveries and the way forward. *Frontiers in Marine Science*, 6.
- Walesby, K., Vialard, J., Minnett, P., Callaghan, A., and Ward, B. (2015). Observations indicative of rain-induced double diffusion in the ocean surface boundary layer. *Geophysical Research Letters*, 42(10):3963–3972.

- Wang, X., Chen, R., Li, C., Chen, Z., Hui, F., and Cheng, X. (2022). An intercomparison of satellite derived Arctic sea ice motion products. *Remote Sensing*, 14(5).
- Webster, P. J., Clayson, C. A., and Curry, J. A. (1996). Clouds, radiation, and the diurnal cycle of sea surface temperature in the tropical western Pacific. *Journal of Climate*, 9(8):1712–1730.
- Weiss, J. (2013). *Drift, Deformation, and Fracture of Sea Ice: A Perspective Across Scales*. Springer Dordrecht.
- Wijesekera, H. and Boyd, T. J. (2001). Upper ocean heat and freshwater budgets. *Encyclopedia of Ocean Sciences*, pages 3079–3083.
- Wijesekera, H., Pauson, C., and Huyer, A. (1999). The effect of rainfall on the surface layer during a westerly wind burst in the western equatorial Pacific. *Journal of Physical Oceanography*, 29(4):612–632.
- Williams, P., Guilyardi, E., Sutton, R., Gregory, J., and Madec, G. (2006). On the climate response of the low-latitude Pacific Ocean to changes in the global freshwater cycle. *Climate Dynamics*, 27.
- Wong, A. P. S., Wijffels, S. E., Riser, S. C., Pouliquen, S., Hosoda, S., Roemmich, D., Gilson, J., Johnson, G. C., Martini, K., Murphy, D. J., Scanderbeg, M., Bhaskar, T. V. S. U., Buck, J. J. H., Merceur, F., Carval, T., Maze, G., Cabanes, C., André, X., Poffa, N., Yashayaev, I., Barker, P. M., Guinehut, S., Belbéoch, M., Ignaszewski, M., Baringer, M. O., Schmid, C., Lyman, J. M., McTaggart, K. E., Purkey, S. G., Zilberman, N., Alkire, M. B., Swift, D., Owens, W. B., Jayne, S. R., Hersh, C., Robbins, P., West-Mack, D., Bahr, F., Yoshida, S., Sutton, P. J. H., Cancouët, R., Coatanoan, C., Dobbler, D., Juan, A. G., Gourrion, J., Kolodziejczyk, N., Bernard, V., Bourlès, B., Claustre, H., D’Ortenzio, F., Le Reste, S., Le Traon, P.-Y., Rannou, J.-P., Saout-Grit, C., Speich, S., Thierry, V., Verbrugge, N., Angel-Benavides, I. M., Klein, B., Notarstefano, G., Poulain, P.-M., Vélez-Belchí, P., Suga, T., Ando, K., Iwasaka, N., Kobayashi, T., Masuda, S., Oka, E., Sato, K., Nakamura, T., Sato, K., Takatsuki, Y., Yoshida, T., Cowley, R., Lovell, J. L., Oke, P. R., van Wijk, E. M., Carse, F., Donnelly, M., Gould, W. J., Gowers, K., King, B. A., Loch, S. G., Mowat, M., Turton, J., Rama Rao, E. P., Ravichandran, M., Freeland, H. J., Gaboury, I., Gilbert, D., Greenan, B. J. W., Ouellet, M., Ross, T., Tran, A., Dong, M., Liu, Z., Xu, J., Kang, K., Jo, H., Kim, S.-D., and Park, H.-M. (2020). Argo data 1999–2019: Two million temperature-salinity profiles and subsurface velocity observations from a global array of profiling floats. *Frontiers in Marine Science*, 7.
- Wunsch, C. (2006). *Discrete Inverse and State Estimation Problems: With Geophysical Fluid Applications*. Cambridge University Press.
- Yamaguchi, R. and Suga, T. (2019). Trend and variability in global upper-ocean stratification since the 1960s. *Journal of Geophysical Research: Oceans*, 124(12):8933–8948.

- Yu, L. (2011). A global relationship between the ocean water cycle and near-surface salinity. *Journal of Geophysical Research*, 116(C10025).
- Yu, L., Josey, S. A., Bingham, F. M., and Lee, T. (2020). Intensification of the global water cycle and evidence from ocean salinity: a synthesis review. *ANNALS of The New York Academy of Sciences*, 1472(1).
- Zhai, J. and Bitz, C. M. (2021). A machine learning model of Arctic sea ice motions.
- Zhang, J., Lindsay, R., Schweiger, A., and Rigor, I. (2012). Recent changes in the dynamic properties of declining arctic sea ice: A model study. *Geophysical Research Letters*, 39(20).
- Zhang, J., Rothrock, D., and Steele, M. (2000). Recent changes in arctic sea ice: The interplay between ice dynamics and thermodynamics. *Journal of Climate*, 13(17):3099 – 3114.
- Zhou, Y. and Newell, R. E. (1998). A proposed algorithm for moisture fluxes from atmospheric rivers. *Monthly Weather Review*, 126(3):725–735.