WIFF1.0: A hybrid machine-learning-based parameterization of Wave-Induced sea-ice Floe Fracture.

Christopher Horvat[a] and Lettie A. Roach[b]

[a]Institute at Brown for Environment and Society, Brown University, Providence, RI, USA
[b]Department of Atmospheric Sciences, University of Washington, Seattle, WA, USA

Correspondence: Christopher Horvat (horvat@brown.edu)

Abstract. Ocean surface waves play an important role in maintaining the marginal ice zone, a heterogenous region occupied by sea ice floes with variable horizontal sizes. The location, width, and evolution of the marginal ice zone is determined by the mutual interaction of ocean waves and floes, as waves propagate into the ice, bend it, and fracture it. In previous work, we developed a one-dimensional "superparameterized" scheme to simulate the interaction between the stochastic ocean surface wave field and sea ice. As this method is computationally expensive and not bitwise reproducible, here we use a pair of neural networks to accelerate this parameterization, delivering an adaptable, computationally-inexpensive, reproducible approach for simulating stochastic wave-ice interactions. Implemented in the sea ice model CICE, this accelerated code reproduces global statistics resulting from the full wave fracture code without increasing computational overheads. The combined model, Wave-Induced Floe Fracture (WIFF v1.0) is publicly available and may be incorporated into climate models that seek to represent the effect of waves fracturing sea ice.

1 Introduction

Sea ice is a composite of individual pieces, called floes. Floes are discrete solid pieces of sea ice with horizontal geometric length scales (sizes) ranging from meters to tens of kilometers. Yet in climate model simulations, sea ice is treated as a continuum field (Hibler, 1979; Golden et al., 2020). To represent the effect of changes to sea ice floe geometry in climate models, modeling centers have started to incorporate prognostic sub-grid-scale parameterizations of the sea ice floe size distribution (FSD) as a core aspect of their model physics (Horvat and Tziperman, 2015; Zhang et al., 2015; Roach et al., 2018; Boutin et al., 2018; Bateson et al., 2020; Hunke et al., 2019). The FSD is a probability distribution $f(r)$, where $f(r)dr$ is the percentage of a grid area comprised of floes with a size between r and $r + dr$.

One major advantage of simulating a sub-grid-scale FSD is it permits the representation of coupled interactions between sea ice floes and ocean surface waves. Ocean surface waves propagate into, are attenuated by, and fracture sea ice in regions known as the marginal ice zone (MIZ, Wadhams et al., 1988; Langhorne et al., 1998; Squire et al., 1995). Wave-affected sea-ice-covered regions are observed to be several million square kilometers in size in both hemispheres (Horvat et al., 2020). When waves propagate into and fracture sea ice, there is no direct change to sea ice concentration or thickness. Instead, this fracture process alters sea ice floe sizes, increasing the sensitivity of the sea ice to external thermodynamic or dynamic forcing

25 (Steele, 1992; Feltham et al., 2006; Horvat et al., 2016; Horvat and Tziperman, 2017) and altering the attenuation of ocean wave energy (Perrie and Hu, 1996; Meylan et al., 2021). Thus there is a hypothesized coupling between sea ice loss and wave activity - waves propagate into the ice, fracturing it, causing it to melt, and enhancing wave propagation (Kohout et al., 2011; Asplin et al., 2012, 2014). Floe fracture by waves is the dominant process driving changes in floe perimeter in summer months when the most lateral melting occurs (Roach et al., 2019).

30 In a climate model, coupled wave-ice feedbacks are related to two sub-grid-scale distributions: the FSD, $f(r)$, and the ocean surface wave spectrum, $S(\lambda)$, where λ is the wavelength and $\int S(\lambda)d\lambda = E$ is the total wave energy. Heuristic parameterizations have been developed to relate bulk properties of the ocean surface wave field to the FSD - generally assuming that fractured floes follow a power-law distribution (Williams et al., 2013; Zhang et al., 2015; Bateson et al., 2020). Yet there is conflicting evidence about whether power-law FSDs are observed in nature (Herman, 2013; Stern et al., 2018a, b; Horvat et al.,

35 2019), and whether power-law size distributions are generated by the process of wave-induced floe fracture (Horvat et al., 2016; Herman et al., 2021).

One challenge in relating $S(\lambda)$ and $f(r)$ is that the response of sea ice is determined by the two-dimensional ocean height field, a stochastic (i.e. one of many possible) representation of the ocean wave spectrum. An approach designed by Horvat and Tziperman (2015) was to generate a high-resolution, one-dimensional ocean surface wave field in every ice-covered grid

40 cell, and explicitly resolve the strain experienced by the ice, the attenuation of wave energy, and the resulting statistics of sea ice fracture. This method (hereafter called SP-WIFF) is broadly analogous to a "super-parameterization" (SP) approach used in coupled climate models for parameterization of cloud effects (Randall et al., 2003; Grabowski, 2004) or oceanic deep convection (Campin et al., 2011). SP-WIFF is included as a component of the sea ice model CICE (Roach et al., 2018; Hunke et al., 2019), and has been coupled to an ocean surface wave model (Roach et al., 2019). Running SP-WIFF incurs high

45 computational costs - in the case of its incorporation in CICE, it overall computation times by an order of magnitude (see Sec. 4). The stochastic generation of wave heights also means SP-WIFF is not bitwise reproducible, often a necessary feature for climate model development.

To provide a computationally inexpensive, flexible, and bit-wise reproducible sub-grid-scale parameterization of wave-induced ice fracture, we here present an accelerated parameterization of wave-induced floe fracture using a pair of neural

50 networks for input classification and sea ice fracture (NN-WIFF). The full code, WIFF1.0 (Wave-Induced Floe Fracture), is publicly available, and contains SP-WIFF and NN-WIFF, along with code for retraining or adapting NN-WIFF to pre-scribed error thresholds, variable input/output variables, and model configurations. Trained using input from coupled CICE-WAVEWATCH 3 simulations, NN-WIFF produces global sea ice variability that is not statistically significantly different from those produced by SP-WIFF, while reducing its computational overhead by more than 95%.

55 2 SP-WIFF: a super-parameterized wave fracture scheme

The super-parameterization of sea ice fracture by ocean waves was derived in Horvat and Tziperman (2015) (Sec. 2.3), its statistical properties explored in Horvat and Tziperman (2017) (Sec. 3.4), was implemented with offline wave-ice interactions

in a sea ice model in Roach et al. (2018) (Sec. 2.4), and was introduced into a fully-coupled wave-ice model in Roach et al. (2019). We summarize it here.

Consider a region corresponding to a climate model grid cell where ocean surface wave energetics are described by a uni-directional wave spectrum $S(\lambda_i)d\lambda_i$ (integrated over all angles). For floes with horizontal size between r and $r + dr$, a fraction of the domain $\Omega(r,t)\,dr\,dt$ (unitless) is fractured by ocean surface waves over a period dt. That fractured area has its own floe size distribution - the fraction of $\Omega dr\,dt$ that now belongs to floes with size between s and $s + ds$ is $F(r,s)\,ds$ (unitless). The rate of change of area of floes of size r due to fracture by ocean surface waves is then,

$$\frac{\partial f(r)}{\partial t} = -\Omega(r,t) + \int\limits_{r}^{\infty} \Omega(s,t)F(s,r)\,ds. \tag{1}$$

The first term in Eq. 1 is the loss of floe area at size category r per unit time, and the second is the increase in floe area in size category r due to the fracture of floes of larger sizes. The limits of integration reflect the fact that any given floe can not fracture into a larger floe, and therefore $F(s,r) = 0$ for $s < r$. Evaluation of Ω and F occurs as follows within each ice thickness category:

S1: The discrete unidirectional wave spectrum $S(\lambda)d\lambda$ is converted to a 1-dimensional ice strain field $\eta(x)$ of length 10 km, (see (Horvat and Tziperman, 2015, eq. 20-21))

S2: A collection of fracture lengths $\{L\}$ is found by finding the distance between successive strain extrema that would fracture a floe. These are binned into a probability distribution $A(r)$, where $A(r)dr$ is the percentage of fracture lengths with length between r and $r + dr$. Note we replace $H(r)$ in Horvat and Tziperman (2015) with $A(r)$ here to avoid confusion with sea ice thickness, H.

S3: **S1-S2** are repeated (increasing the size of $\{L\}$) until convergence, defined as a mean absolute difference between successively updated values of $A(r)$ below 5×10^{-4}.

The histogram $A(r)dr$ is then used to compute Ω and F,

$$\Omega(r) = \int\limits_{0}^{r} A(r)dr / \int\limits_{0}^{\infty} A(r)dr \tag{2}$$

$$F(s,r)ds = A(r)dr / \int\limits_{0}^{s} A(r)dr. \tag{3}$$

SP-WIFF therefore involves a stochastic component (**S1**) and an expensive peak-finding and histogram-generation component (**S2**). Its execution time is uncertain because the number of steps until convergence (**S3**) is not known in advance. Previously, to circumvent these challenges, SP-WIFF was implemented in CICE using a fixed phase in step **S1** and only iterated once in step **S3**. We refer to this single-iteration implementation as SP-WIFF$_1$. Despite this simplification, CICE runtimes are increased by approximately a factor of 4 (see Sec. 4) over benchmark estimates without wave fracture. We seek to improve the performance of SP-WIFF by replacing **S1**-**S3** with a neural network trained on SP-WIFF input and output data, which we call NN-WIFF.

3 Accelerating the parameterization of coupled ice-wave interactions

The NN-WIFF code replaces the full super-parameterization (SP-WIFF) with a parameterization that replicates its statistics
90 but has reduced computational overhead. It consists of two full-connected feedforward neural networks trained using the
Python Keras deep learning library, one 100x100 input classification scheme for determining whether NN-WIFF will run, and
a second with five hidden layers of 100 nodes each for generating fracture histograms. Network architectures were chosen
using a metalearning approach varying loss functions, and numbers of nodes and layers (we provide code for altering network
size in the WIFF1.0 release). Training is performed using the 'adam' optimization scheme using the resilient backpropagation
95 method. Rectified linear unit activation functions are used for each of the hidden neurons, and we use a softmax activation for
final output layer in both cases as we seek a binary value (for the classifier) and a histogram that sums to one (for the fracture
histogram network).

Because SP-WIFF can be run on any arbitrary input, we may generate a training dataset of unlimited size. In this study,
we use data from a single year (2009) of a coupled simulation of waves and sea ice (as in the simulation FSD-WAVEv2 in
100 Roach et al. (2019), but here we use a fully-converged wave fracture parameterization instead of a single application of **S1**-**S2**,
i.e. SP-WIFF instead of SP-WIFF$_1$). We take 6-hourly wave spectra, sea ice concentration, sea ice thickness, and fracture
histograms in both hemispheres, a total of 180 million possible input vectors (most without sea ice). We prune training data by
including only locations where the local sea ice thickness is less than 10 m, sea ice concentration is greater than 0.01, and the
significant wave height $H_s = 4\sqrt{E}$ is greater than 0.1 m. These thresholds are also checked before executing SP-WIFF in the
105 version of WIFF released here. This eliminates spurious wave fracture calls in areas of anomalous sea ice conditions or low
wave energies. We find a total of 17.9 million sea ice points, of which 5.1 million meet the criterion above.

Thus in total, the full training data set comprises $N = 5.1$ million input vectors of 25 spectral amplitudes, 1 ice thickness,
and 1 ice concentration. We assume these include the potential phase space of inputs to the SP-WIFF code, and the potential
range of wave energies and sea ice states. These input vectors are identified with a set of output vectors from the "true" SP-
110 WIFF output - in this case 5.1 million 12-category floe probability distributions $A_i dr_i$, of which 2.6 million are non-zero (i.e.
SP-WIFF leads to floe fracture). We choose to use the 12 category FSD presented in Roach et al. (2019), however this can
readily be altered for different uses. In the WIFF1.0 release, we include code for: (see code and data availability statement)

– Defining a custom training data set using prescribed input spectra and thicknesses.

– Defining a custom output histogram size (i.e. different FSD categories).

115 – Designing and training the input classification scheme.

– Designing and training the histogram generation scheme.

We also provide data that was used to train networks and produce the results shown here (see Code and Data Availability).
We used coupled model output of SP-WIFF for the 12-category floe probability distributions in our training dataset. These
can also be generated offline, and we provide offline code (*SP_WIFF_Standalone.m*) for generating output histograms for any
120 arbitrary input that is bit-for-bit the same as the code in SP-WIFF.
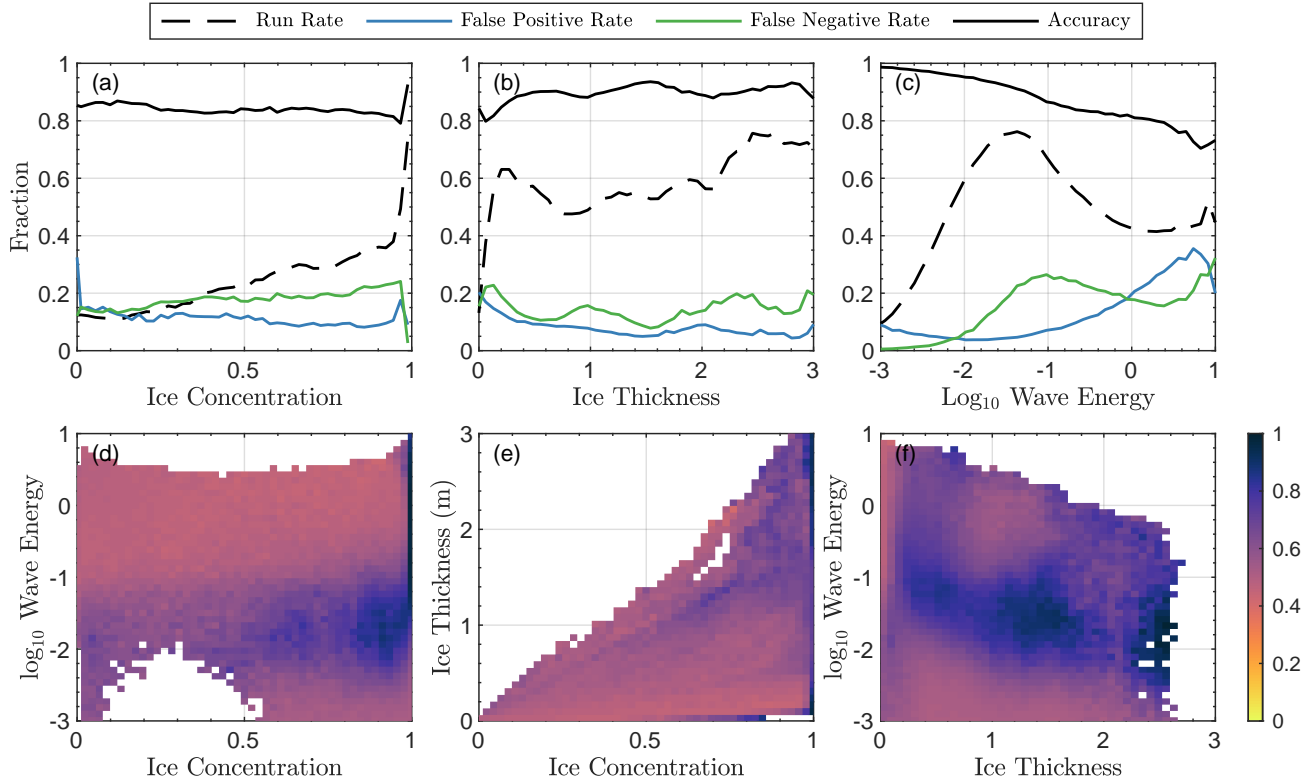
Figure 1. Performance of Input Classification for NN-WIFF (a) Run rate (dashed black), false positive rate (blue), false negative rate (green), and overall accuracy (black) for input classifier as a function of sea ice concentration. (b,c) Same as (a) for (b) sea ice thickness and (c), $\log_{10}$ wave energy. (d-f) Heat map of average classifier accuracy as a function of (d) ice concentration and $\log_{10}$ wave energy, (e) ice concentration and ice thickness, and (f) ice thickness and $\log_{10}$ wave energy.

3.1 Input Classification

The histogram $A_i dr_i$ is a collection of numbers that sums to one, and is the target of our parameterization acceleration. This motivates the use of a softmax activation layer (as the output sums to one). Yet when SP-WIFF is executed, sea ice does not always fracture, and a true output histogram could be a vector of zeros. We therefore require a classification layer that

125 determines when a histogram will be produced by NN-WIFF, and when it will not, to reduce the potential for false positives. We found the bulk of potential false positives could be eliminated by introducing the wave energy cutoff of $H_s = 4\sqrt{E} = 0.1$m above, and therefore the classification network is implemented after that cutoff is applied.

The training data for the classifier is the 5.1 million input vectors with corresponding binary values (fracture/no fracture). This is randomly partitioned 70/30 into training (3.6 million) and validation (1.5 million) datasets, with an even proportion in

130 each set (51%) leading to fracture. The input classifier then takes an input vector and returns a single value, $0 < \chi < 1$, which

executes the NN-WIFF scheme if $\chi < \chi_{crit}$. The classification threshold χ_{crit} is determined by optimizing the error properties of the classifier when evaluated against the validation dataset. We optimize the total error rate, ϵ, where,

$$\epsilon = \frac{\text{False Positives + False Negatives}}{\text{Total Input Vectors}} \tag{4}$$

We found that ϵ is minimized for $\chi_{crit} = 0.54$ at $\epsilon = 12.5\%$. This error rate corresponds to a false positive rate of 10.6% and a

135   false negative rate of 14.0%.

Fig. 1 shows error characteristics of the classifier evaluated on the validation dataset. In Fig. 1(a-c), we show the relationship between false positive rate (blue) and false negative rate (green) as a function of (Fig. 1a) sea ice concentration, (Fig. 1a) sea ice thickness, and (Fig. 1c) $\log_{10} E$. We also plot the run rate (dashed black line), or the fraction of times NN-WIFF is executed, as well as the overall accuracy (black line), equal to 1-ϵ.

140   Generally, the input classifier has high accuracy across the range of ice concentrations and thicknesses, with a flat distribution of false positives and false negatives. Lower accuracy occurs at the highest range of wave energies, where both false positives and false negatives are higher than 20%. We plot two-dimensional maps of the overall accuracy in Fig. 1(d) ($\log_{10} E$ vs ice concentration), Fig. 1(e) (ice thickness vs ice concentration), and Fig. 1(f) ($\log_{10} E$ vs ice thickness). Some regions of phase space have classifier accuracies as low as 42%. The classifier accuracy is high in the regions with the highest run rates - high

145   sea ice concentrations, thicknesses, and wave energies between 0.01 and 0.1 m^2. As we explore below, these regions represent the largest portion of potential input states, and the 12.5% error in classification does not contribute to statistically significant difference in climate model output. Still, improvements in the input classification may be achievable if desired, for example for higher wave energies, and we include code to re-train the classifier in WIFF1.0.

3.2   Neural Network for Generating Fracture Histograms

150   The fracture histogram network takes a length-27 input vector and outputs a vector $\hat{A}_i \, dr_i$ which sums to one and approximates the histogram $A_i \, dr_i$. As a result, sea ice floes will fracture if this input wave and sea ice statistics pass the input classification. To assess its performance, we use a custom loss function that minimizes the "representative size error", which is the absolute difference in the predicted histogram in each floe size category, weighted by the floe size in that category,

$$RSE = \sum_i \left| \hat{A}_i - A_i \right| r_i \, dr_i, \tag{5}$$

155   and has units of meters. Available training data is randomly partitioned into a training dataset of 70% of the input vectors, and a validation dataset of the remaining 30% ($\sim$800,000 input vectors). Training is performed until the validation loss fails to improve over 20 subsequent steps. Indexing training data by j, the overall network performance is the average error for the training data,

$$P = \frac{1}{N} \sum_1^N RSE_j \approx 30\text{m}, \tag{6}$$

160   where N is the total number of inputs used to train the network.
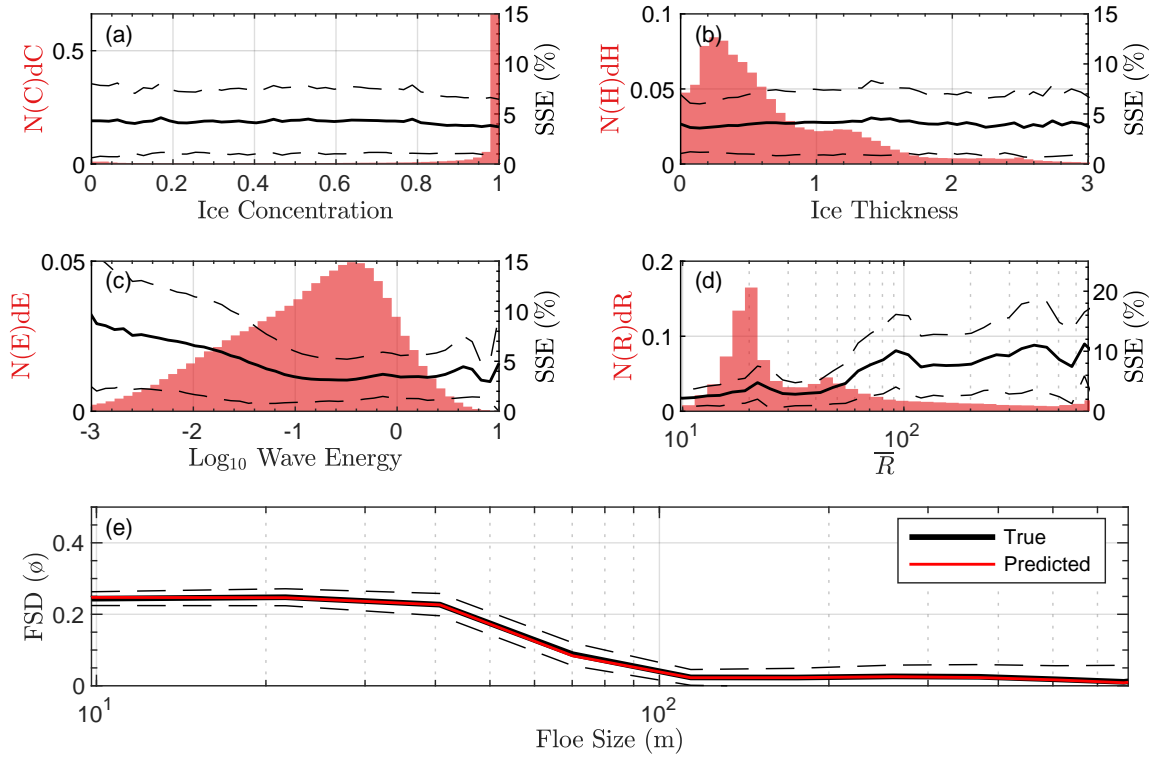
Figure 2. Performance of NN-WIFF on validation data. (a, left axis) Histogram of sea ice concentration, C, in the validation dataset. (a,right axis) Median standardized size error (SSE) as a function of C (solid line). Dashed lines are interquartile distance of SSE. (b-d) Same as (a), but for (b) sea ice thickness (c) $\log_{10}$ wave energy, and (d) representative radius of the "true" output histogram. (e) Mean histogram output $A_i dr_i$ for (black) validation dataset and (red) parameterization output $\hat{F}_i dr_i$. Dashed lines are standard deviation of error between $A_i dr_i$ and $\hat{A}_i dr_i$ in each floe size category.

In Figure 2 we display aggregated fit characteristics for the histogram neural net. Figures 2(a-d) display number histograms (red, left axis) of sea ice concentration (C), ice thickness (H), and $\log_{10}$E from the validation dataset, as well as of the "true" representative radius $\overline{R}$, where

$$\overline{R} = \sum_i r_i A_i dr_i. \tag{7}$$

165 from validation dataset input. To evaluate the standardized neural net error, we also compute the standardized size error (SSE) as,

$$SSE = \frac{RSE}{\overline{R}}, \tag{8}$$
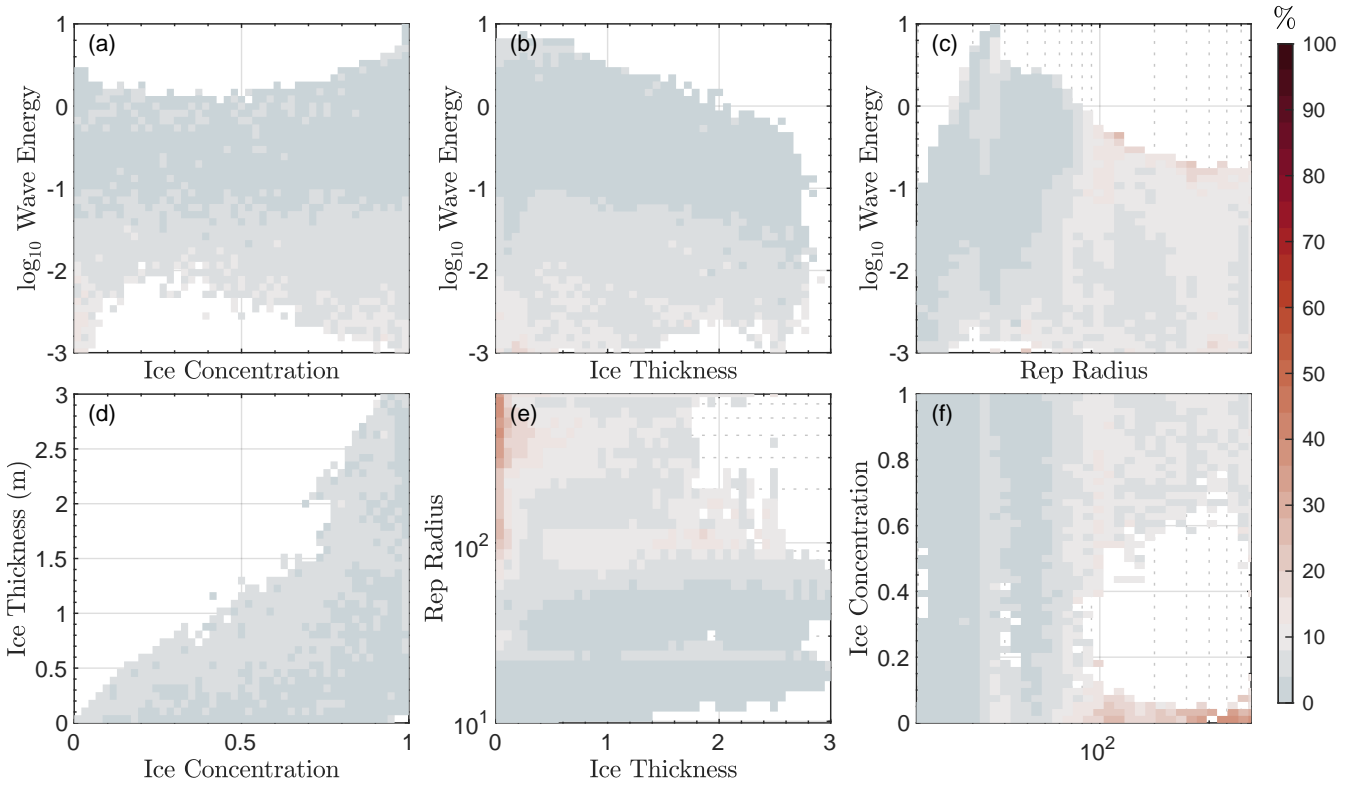
Figure 3. Examination of Training Errors in NN-WIFF. (a) 2-D map of median SSE as a function of (x-axis) sea ice concentration and (y-axis) $\log_{10}$E from validation dataset. (b-f) same as (a) but for (b) H and $\log_{10} E$, (c) $\overline{R}$ and $\log_{10} E$, (d) C and H, (e) H and $\log_{10} E$, and (f) $\overline{R}$ and C. Regions are not colored for less than 0.005% (∼40) points. Colormap white point is at SSE=10%.

which weights errors by the expected mean floe size. Across the entire validation dataset, median SSE is 3.9%.

There is little dependency of SSE on sea ice concentration (2a) or sea ice thickness (2b). Median SSE errors are below
170 5% across all ice concentrations and thicknesses considered here, with the peak of the interquartile range below 10%. While median errors for both wave energy and $\overline{R}$ also lie below 10% for all input wave energies and true output floe sizes, there are higher SSE values for low wave energies and high $\overline{R}$ values. Large values of $\overline{R}$ can occur for a small number of potential floe fractures, which typically occurs for low wave energies that do not repeatedly fracture the sea ice. In general, the neural network (red line, Fig. 2e) reproduces the mean fracture histogram $\overline{A}_i$ (black line, Fig. 2e) accurately at each floe size category.
175 Median error between $\overline{A}_i dr_i$ and $\widehat{A}_i dr_i$ is less than 6.2% in all categories, and across all 12 categories averages 3.3%.

To understand what regions of the phase space give rise to error, in Fig. 3(a-c) we plot 2-dimensional heat maps of SSE, evaluated on the validation dataset, for each of the six combinations of ice thickness, concentration, $\overline{R}$, and log wave energy. Colors are shown only for those coordinates with at least 0.005% of the data (40 points), and the whitepoint of the colormap is set for an error of 10%. Blue values are those with SSE < 10%.

180   Fig. 2 suggests that the input vectors with largest error are those with low input wave energy and high resulting $\overline{R}$. In general, we find similar results when examining these two-dimensional maps. In general, we find low error (SSE<=5%) across the range of input sea ice thickness and concentration (Fig. 3a). However expanding the error map in $\overline{R}$ as a function of thickness (Fig. 3b), concentration (Fig. 3f), and wave energy (Fig. 3c) reveals that the largest part of training error comes from regions with low input sea ice concentration, small thickness, and low wave energy. This combination of inputs gives rise to

185   high output $\overline{R}$ values, which have the highest SSE values. Although locations with low sea ice concentration are not used as WIFF input, we do not presently exclude locations with low sea ice thickness - that these regions produce pathological output suggests they might be isolated or excluded from future versions of NN-WIFF.

4   Sea Ice Model Results

We next examine the relationship between versions of the WIFF code when implemented in a free-running sea ice model.

190   Here we use output from the global sea ice model CICE6.1.4 (Hunke et al., 2019) in stand-alone mode forced by the JRA55 atmospheric reanalysis (JRA-55 and Japan Meteorological Agency, 2013) at a nominal 1^o spatial resolution. We include time-varying ocean surface wave spectra as sea ice model forcing, taken from output of the CICE-Wavewatch III coupled integration described in Roach et al. (2019). The model is run over the year 2005 to provide different boundary conditions and initialization from the 2009 data used as input to the network training.

195   We examine results from three separate methods for obtaining the output histogram $A_i dr_i$:

1. (SP-WIFF) using the full SP-WIFF model but the converged solution described in Sec. 2

2. (SP-WIFF$_1$) the SP-WIFF model but using a single iteration of steps **S1**-**S2** with fixed phase.

3. (NN-WIFF) the trained neural network described above.

The histogram output $A_i dr_i$ from CICE is a daily average over 24 hour timesteps. When wave fracture does not occur, $A_i dr_i$

200   values are set to zero. Thus though $A_i dr_i$ is normalized to one if wave fracture occurs, if a grid cell does not always fracture during one of the averaging timesteps, the CICE history output will not be normalized to one. We choose to include such points in this analysis, renormalizing histogram outputs to sum to one.

Over the course of this year, we find that in the NN-WIFF simulations, there are 686,000 gridpoint-days during which NN-WIFF is executed, for a total of 12.3 million NN-WIFF calls overall. This compares to 746,000 and 14.1 million for SP-

205   WIFF and SP-WIFF-1. The discrepancy between function calls is due to the use of the input classifier, which introduces false negatives as discussed in Sec. 3.1.

In Fig. 4(a,b) we repeat Fig. 2(e) but now examining (a) Arctic and (b) Antarctic averages of $A_i dr_i$ values for each of the three methods. To estimate differences in fracturing, Fig. 4 shows data only from locations and times where WIFF leads to sea ice fracture in all three of the simulations, and where the daily-average wave energy exceeds 0.1 m^2. Dashed lines show the

210   interquartile range of SP-WIFF values. Median values are similar between all three implementations.
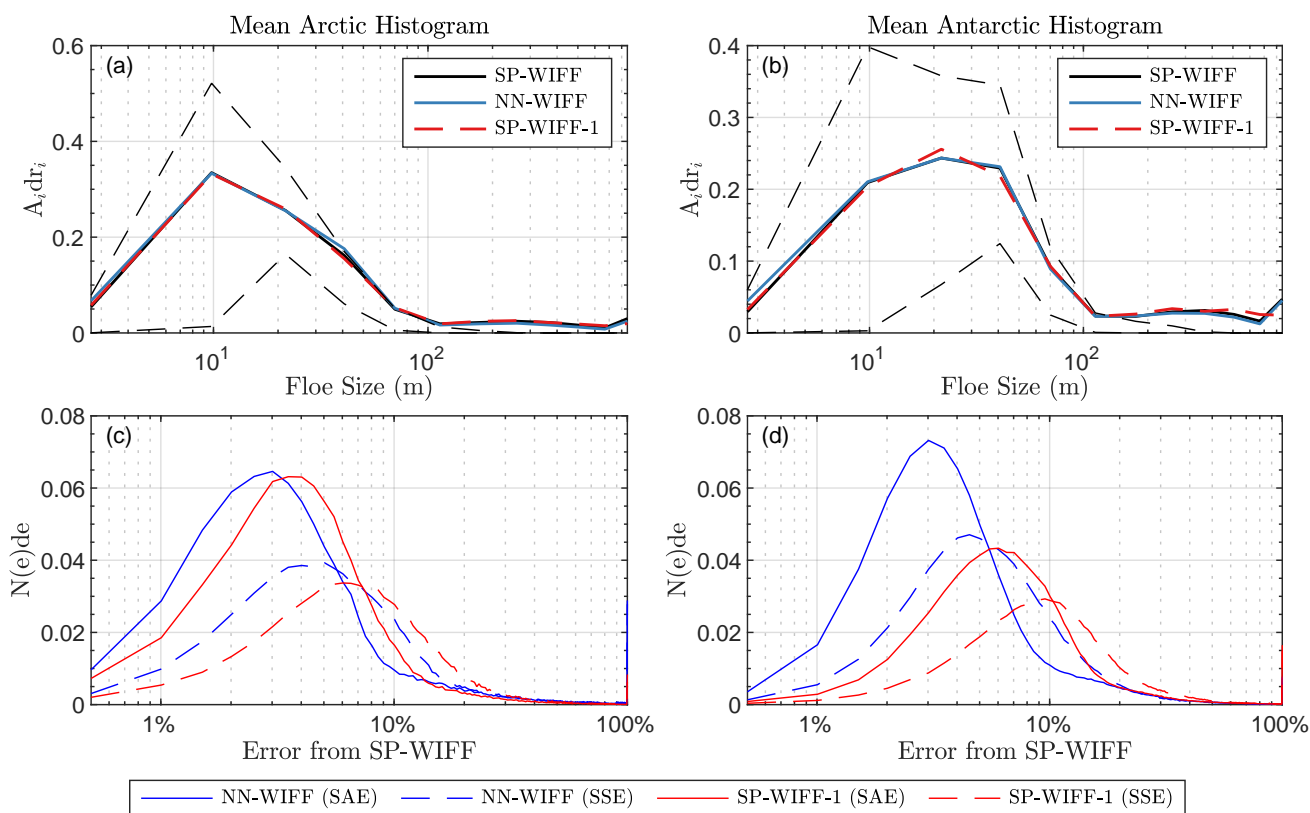
Figure 4. WIFF error characteristics in a climate model. (a) Mean Arctic histogram output $A_i dr_i$ in free-running climate simulation using (black) SP-WIFF, (blue) NN-WIFF, and (red dashed) SP-WIFF$_1$. Space/time points included in histogram if all three cause wave fracture. Interquartile range of SP-WIFF outputs given as dashed black lines. (b) Same as (a) for the Antarctic. (c) Histogram of errors between SP-WIFF and (blue) NN-WIFF or (red) SP-WIFF-1 in the Arctic. Solid lines are sum of absolute error (SAE). Dashed lines are standardized size error (SSE). (d) Same as (c), for the Antarctic.

We plot histograms of the errors between SP-WIFF and SP-WIFF$_1$ (red) and NN-WIFF (blue) in both hemispheres as Fig. 4(c-d) from the same sample used in producing Fig. 4(a,b). Dashed lines are SSE and solid lines are sum of absolute error (SAE) in each category. In both metrics, and in both hemispheres, NN-WIFF outperforms SP-WIFF$_1$ when it is executed. In the Arctic, median SSE is 8.9% (NN-WIFF) vs 10.5% (SP-WIFF$_1$), and median SAE is 5.8% (NN-WIFF) vs 5.9% (SP-WIFF$_1$).
215   In the Antarctic, median SSEs are 8.1% (NN-WIFF) vs 13.4% (SP-WIFF$_1$) and median SAEs are 5.52% (NN-WIFF) vs 9.02% (SP-WIFF$_1$).

In Fig. 5(top row) we show monthly-average differences in sea ice concentration between SP-WIFF and NN-WIFF in March and September at each hemisphere. There is little difference between the two, with a maximum global area difference of 47,000 km^2 in January (0.3% of total sea ice area). We repeat this analysis on sea ice volume per unit area in Fig. 5(middle row). We
220   again find little difference throughout the year, with a global maximum volume difference of 33 km^3 in January (0.12% of total sea ice volume). We also examine the log ratio of SP-WIFF to NN-WIFF mean floe size (representative radius) as the bottom row in Fig. 5. Some areas have larger floe sizes in SP-WIFF compared to NN-WIFF - those where the discrepancy in mean floe sizes is a multiple of two or more, i.e. $|\log_{10}(R_{NN}/R_{SP})| > \log_{10}(2)$ have a maximum global area of 1100 km^2 in April, or 4.8% of total sea ice area. These regions are visible as locations lying at the boundary between the MIZ and sea ice pack,
225   i.e. locations where NN-WIFF does not trigger wave fracture but SP-WIFF does. We evaluate these differences between model integrations using the CICE quality-control check (Roberts et al., 2018). Although there is not bit-for-bit reproducibility, we find that the different implementations of WIFF are not "climate changing" according to the two-stage paired thickness test. This demonstrates that differences in WIFF implementation do not have an emergent effect on sea ice model state.

NN-WIFF substantially accelerates the parameterization of wave-induced floe fracture. Integrating these simulations using
230   on the Cheyenne supercomputer with daily I/O on a 384x320 displaced pole grid, using 1 node and 24 cores, the standard SP-WIFF executes 1 year in approximately 24 hours. The single-iteration version SP-WIFF-1 takes 3h30. The NN-WIFF model takes 1h10 minutes. For comparison, a standard CICE integration with the FSD but no waves takes 56 minutes, and a CICE integration with no FSD and no waves takes 45 minutes. (Note that these runtimes do not include integration of a wave model, as this was run offline and provided as a forcing to the sea ice model.) Thus while execution times will vary with computer
235   architecture and user-specified compiler settings, NN-WIFF reduces the overhead associated with simulating wave-ice fracture without significant added computational cost.

5   Conclusions

Here we have presented WIFF1.0, code for performing wave-induced floe fracture in sea ice models that are coupled to an ocean surface wave model. While the full "super-parameterized" version of WIFF (SP-WIFF) is capable of simulating wave-
240   affected marginal ice zones with active wave-ice interactions, it results in high computational expense which renders it difficult to use in long climate integrations. Instead, we trained a pair of neural networks (NN-WIFF) to accelerate this parameterization using a large set of true climate model input/output pairs, which results in a median "standardized size error" of 3.9% when evaluated on a set of 800,000 input states, and 8.4% in free-running climate model simulations.

Because of the use of a classification layer, NN-WIFF does introduce some false positives and negatives - with a classifier
245 accuracy of 87.5%. Yet the "overall" accuracy of calls to NN-WIFF is 96.5%. Because we now employ an energy threshold
for determining whether NN-WIFF would be called, the classifier network is only called on a reduced subset of sea ice points.
The energy threshold eliminates potential calls to NN-WIFF for approximately 72% of total sea ice points, thus the likelihood
that NN-WIFF is inappropriately called at any given sea ice point is $0.125 \times 0.28 = 3.5\%$.

Compared to a single-iteration version of SP-WIFF that was previously employed for computational reasons, NN-WIFF
250 produces improved error characteristics, notably reduced sum and "representative" absolute errors, and is significantly less
expensive. Global patterns of sea ice area and volume are statistically indistinguishable between NN-WIFF throughout a year
of climate model simulations, with minor differences in representative floe sizes at the boundary between the MIZ and pack
ice.

The network has been trained on present-day (2009) sea ice conditions from a single model, and therefore may have less
255 success for different climates where the phase space of wave spectra and ice thicknesses changes. We attempted to use a broad
range of input model vectors that represent a full seasonal cycle of sea ice and wave statistics in both hemispheres. The appeal
of parameterization acceleration is that new versions of NN-WIFF can be trained by directly using other climate model output,
and model input variables and thresholds can be adjusted to provide better fits to this model. We provide code in the WIFF1.0
repository for running SP-WIFF in a standalone fashion, collating training data from a climate model, and (re)-training the
260 network, as well as producing the plots used in this paper.

Because of the ease of obtaining training data from climate model output, this parameterization acceleration approach has,
and could continue to, find a use in many other aspects of climate model simulations (Rasp et al., 2018), for similarly expensive
parameterizations, especially those that are not the solution to primitive equations Pal et al. (2019); Brenowitz et al. (2020).
For example, sea ice rheological parameterizations generally rely on complex implicit solvers run to a specified tolerance, but
265 may, much like SP-WIFF, be able to be decomposed into a series of neural-net "black boxes".

Our primary goal in developing NN-WIFF is to provide a method for accelerating the existing SP-WIFF paramterization,
which will permit the implementation of wave-ice fracture parameterizations in climate-scale sea ice models. In general, NN-
WIFF allows for a cost-effective implementation of wave-ice fracture in sea ice models with no significant increase in runtime.
Because a neural network is a simple set of elementary functions, it can be employed straightforwardly in any climate model
270 that desires to simulate the fracture of sea ice by ocean surface waves.

Code and data availability. Code for CICE6 and Icepack is available at https://github.com/lettie-roach/CICE/tree/mlwave and https://github.
com/lettie-roach/Icepack/tree/mlwave, and will be given a public DOI upon paper acceptance and incorporation into the core CICE code.
Code for developing training data and training the neural networks is available on Github at https://github.com/chhorvat/WIFF-Model and
is archived at doi.org/10.5281/zenodo.5171498. Model output (from the three 2005-forced WIFF configurations) and training data (from the
275 2009 SP-WIFF simulation) are available on Zenodo at doi.org/10.5281/zenodo.5106703 and doi.org/10.5281/zenodo.5108636, respectively.
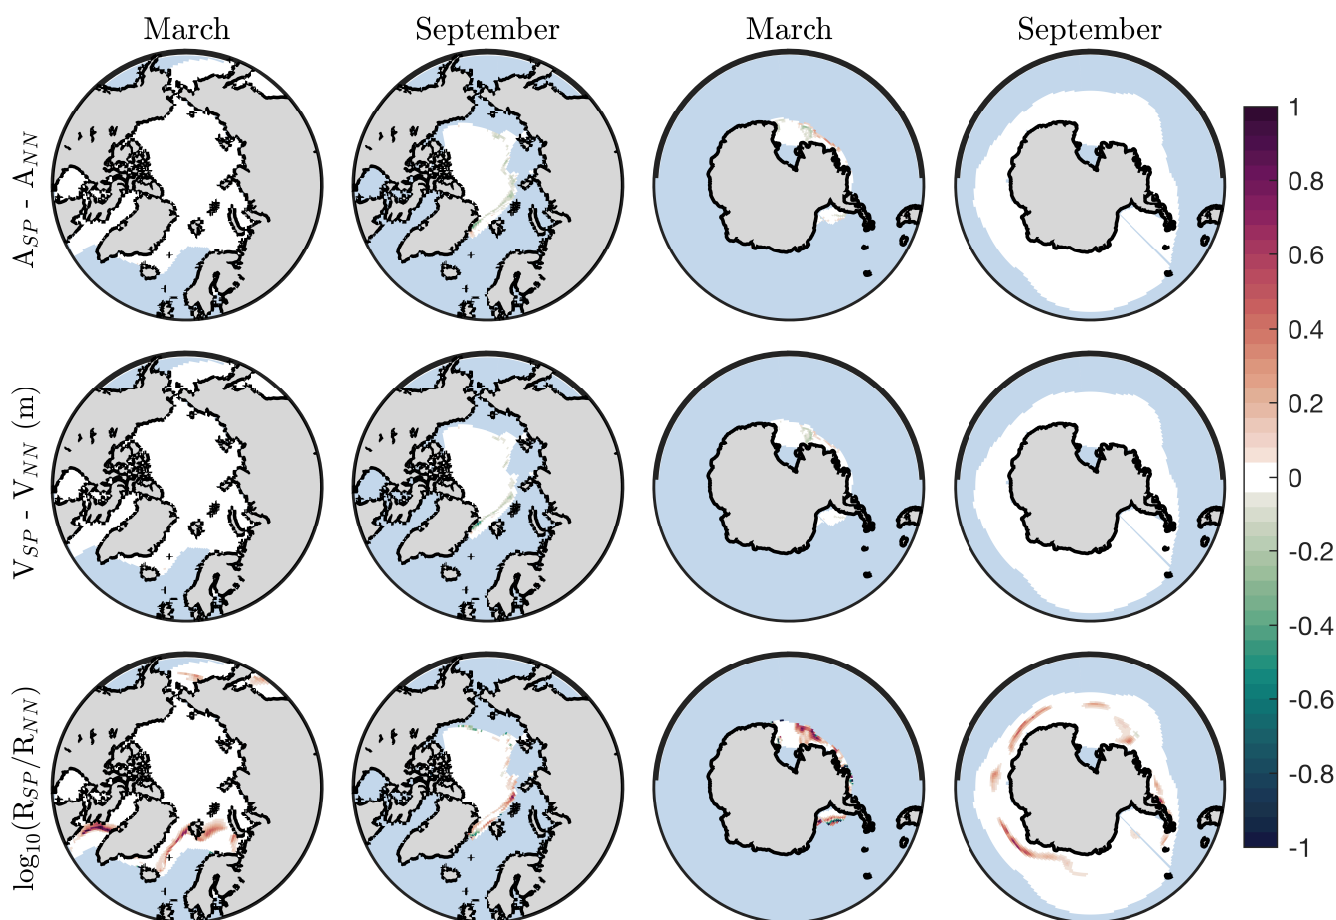
Figure 5. Implementation of WIFF in a climate model. (Top row) Difference in monthly sea ice area, A, between $SP - WIFF$ and $NN - WIFF$ in March and September at each hemisphere. (Middle row) Same as top row, for sea ice volume, V. (Bottom row) Same, for $\log_{10}$ of ratio of representative radius.

Author contributions. Both authors designed the study. CH trained and validated the neural networks and wrote the manuscript. LR incorporated NN-WIFF into CICE and performed model runs for training and validation.

Competing interests. The authors declare no competing interests

Acknowledgements. CH was supported by NASA grant 80NSSC20K0959 and by Schmidt Futures – a philanthropic initiative that seeks to
280 improve societal outcomes through the development of emerging science and technologies. CH thanks the National Institute of Water and
Atmospheric Research in Wellington, NZ for their hospitality during this work. LR was supported by New Zealand Ministry of Business,
Innovation and Employment under Science Investment Contract C01X1914 and US National Science Foundation grant OPP-1643431.

References

Asplin, M. G., Galley, R., Barber, D. G., and Prinsenberg, S.: Fracture of summer perennial sea ice by ocean swell as a result of Arctic storms, Journal of Geophysical Research: Oceans, 117, 1–12, https://doi.org/10.1029/2011JC007221, 2012.

Asplin, M. G., Scharien, R., Else, B., Howell, S., Barber, D. G., Papakyriakou, T., and Prinsenberg, S.: Implications of fractured Arctic perennial ice cover on thermodynamic and dynamic sea ice processes, Journal of Geophysical Research: Oceans, 119, 2327–2343, https://doi.org/10.1002/2013JC009557, 2014.

Bateson, A. W., Feltham, D. L., Schröder, D., Hosekova, L., Ridley, J. K., and Aksenov, Y.: Impact of sea ice floe size distribution on seasonal fragmentation and melt of Arctic sea ice, The Cryosphere, 14, 403–428, https://doi.org/10.5194/tc-14-403-2020, 2020.

Boutin, G., Ardhuin, F., Dumont, D., Sévigny, C., Girard-Ardhuin, F., and Accensi, M.: Floe size effect on wave-ice interactions: possible effects, implementation in wave model, and evaluation, Journal of Geophysical Research: Oceans, 123, 4779–4805, https://doi.org/10.1029/2017JC013622, 2018.

Brenowitz, N. D., Beucler, T., Pritchard, M., and Bretherton, C. S.: Interpreting and stabilizing machine-learning parametrizations of convection, Journal of the Atmospheric Sciences, 77, 4357–4375, https://doi.org/10.1175/JAS-D-20-0082.1, 2020.

Campin, J.-M., Hill, C., Jones, H., and Marshall, J.: Super-parameterization in ocean modeling: Application to deep convection, Ocean Modelling, 36, 90–101, https://doi.org/10.1016/j.ocemod.2010.10.003, 2011.

Feltham, D. L., Untersteiner, N., Wettlaufer, J. S., and Worster, M. G.: Sea ice is a mushy layer, Geophysical Research Letters, 33, L14 501, https://doi.org/10.1029/2006GL026290, 2006.

Golden, K. M., Bennetts, L. G., Cherkaev, E., Eisenman, I., Feltham, D., Horvat, C., Hunke, E., Jones, C., Perovich, D. K., Ponte-Castaneda, P., Strong, C., Sulsky, D., and Wells, A. J.: Modeling Sea Ice, Notices of the American Mathematical Society, 67, 1, https://doi.org/10.1090/noti2171, 2020.

Grabowski, W. W.: An Improved Framework for Superparameterization, Journal of the Atmospheric Sciences, 61, 1940–1952, https://doi.org/10.1175/1520-0469(2004)061<1940:AIFFS>2.0.CO;2, 2004.

Herman, A.: Numerical modeling of force and contact networks in fragmented sea ice, Annals of Glaciology, 54, 114–120, https://doi.org/10.3189/2013AoG62A055, 2013.

Herman, A., Wenta, M., and Cheng, S.: Sizes and shapes of sea ice floes broken by waves - a case study from the East Antarctic coast, Frontiers in Earth Science, 9, 390, https://doi.org/10.3389/feart.2021.655977, 2021.

Hibler, W. D.: A Dynamic Thermodynamic Sea Ice Model, J. Phys. Oceangr., 9, 815–846, https://doi.org/10.1175/1520-0485(1979)009<0815:ADTSIM>2.0.CO;2, 1979.

Horvat, C. and Tziperman, E.: A prognostic model of the sea-ice floe size and thickness distribution, Cryosphere, 9, 2119–2134, https://doi.org/10.5194/tc-9-2119-2015, 2015.

Horvat, C. and Tziperman, E.: The evolution of scaling laws in the sea ice floe size distribution, Journal of Geophysical Research: Oceans, 122, 7630–7650, https://doi.org/10.1002/2016JC012573, 2017.

Horvat, C., Tziperman, E., and Campin, J.-M.: Interaction of sea ice floe size, ocean eddies, and sea ice melting, Geophysical Research Letters, 43, 8083–8090, https://doi.org/10.1002/2016GL069742, 2016.

Horvat, C., Roach, L. A., Tilling, R., Bitz, C. M., Fox-Kemper, B., Guider, C., Hill, K., Ridout, A., and Shepherd, A.: Estimating the sea ice floe size distribution using satellite altimetry: theory, climatology, and model comparison, The Cryosphere, 13, 2869–2885, https://doi.org/10.5194/tc-13-2869-2019, 2019.

320   Horvat, C., Blanchard-Wrigglesworth, E., and Petty, A.: Observing waves in sea ice with ICESat-2, Geophysical Research Letters, https://doi.org/10.1029/2020GL087629, 2020.

Hunke, E., Allard, R., Bailey, D. A., Blain, P., Bouchat, A., Craig, T., Dupont, F., DuVivier, A., Grumbine, R., Hebert, D. A., Holland, M., Jeffery, N., Lemieux, J.-F., Rasmussen, T., Ribergaard, M., Roberts, A., Turner, M., and Winton, M.: CICE-Consortium/CICE: CICE Version 6.0.2, Zenodo, https://doi.org/10.5281/zenodo.3516944, 2019.

325   JRA-55 and Japan Meteorological Agency: JRA-55: Japanese 55-year Reanalysis, Daily 3-Hourly and 6-Hourly Data, Research Data Archive at the National Center for Atmospheric Research, Computational and Information Systems Laboratory, https://doi.org/10.5065/D6HH6H41, 2013.

Kohout, A. L., Meylan, M. H., and Plew, D. R.: Wave attenuation in a marginal ice zone due to the bottom roughness of ice floes, Annals of Glaciology, 52, 118–122, https://doi.org/10.3189/172756411795931525, 2011.

330   Langhorne, P. J., Squire, V. A., Fox, C., and Haskell, T. G.: Break-up of sea ice by ocean waves, Annals of Glaciology, 27, 438–442, https://doi.org/10.3189/S0260305500017869, 1998.

Meylan, M. H., Horvat, C., Bitz, C. M., and Bennetts, L. G.: A floe size dependent scattering model in two- and three-dimensions for wave attenuation by ice floes, Ocean Modelling, 161, 101 779, https://doi.org/10.1016/j.ocemod.2021.101779, 2021.

Pal, A., Mahajan, S., and Norman, M. R.: Using Deep Neural Networks as Cost-Effective Surrogate Models for Super-Parameterized E3SM
335   Radiative Transfer, Geophysical Research Letters, 46, 6069–6079, https://doi.org/10.1029/2018GL081646, 2019.

Perrie, W. and Hu, Y.: Air–Ice–Ocean Momentum Exchange. Part 1:Energy Transfer between Waves and Ice Floes, Journal of Physical Oceanography, 26, 1705–1720, https://doi.org/10.1175/1520-0485(1996)026<1705:AMEPTB>2.0.CO;2, 1996.

Randall, D., Khairoutdinov, M., Arakawa, A., and Grabowski, W.: Breaking the Cloud Parameterization Deadlock, Bulletin of the American Meteorological Society, 84, 1547–1564, https://doi.org/10.1175/BAMS-84-11-1547, 2003.

340   Rasp, S., Pritchard, M. S., and Gentine, P.: Deep learning to represent subgrid processes in climate models, Proceedings of the National Academy of Sciences of the United States of America, 115, 9684–9689, https://doi.org/10.1073/pnas.1810286115, 2018.

Roach, L. A., Horvat, C., Dean, S. M., and Bitz, C. M.: An emergent sea ice floe size distribution in a global coupled ocean–sea ice model, Journal of Geophysical Research: Oceans, 123, 4322–4337, https://doi.org/10.1029/2017JC013692, 2018.

Roach, L. A., Bitz, C. M., Horvat, C., and Dean, S. M.: Advances in Modeling Interactions Between Sea Ice and Ocean Surface Waves,
345   Journal of Advances in Modeling Earth Systems, 11, 4167–4181, https://doi.org/10.1029/2019MS001836, 2019.

Roberts, A. F., Hunke, E. C., Allard, R., Bailey, D. A., Craig, A. P., Lemieux, J.-F., and Turner, M. D.: Quality control for community-based sea-ice model development, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 376, 20170 344, https://doi.org/10.1098/rsta.2017.0344, 2018.

Squire, V. A., Dugan, J. P., Wadhams, P., Rottier, P. J., and Liu, a. K.: Of Ocean Waves and Sea Ice, Annual Review of Fluid Mechanics, 27,
350   115–168, https://doi.org/10.1146/annurev.fl.27.010195.000555, 1995.

Steele, M.: Sea ice melting and floe geometry in a simple ice-ocean model, Journal of Geophysical Research: Oceans, 97, 17 729–17 738, https://doi.org/10.1029/92JC01755, 1992.

Stern, H. L., Schweiger, A. J., Stark, M., Zhang, J., Steele, M., and Hwang, B.: Seasonal evolution of the sea-ice floe size distribution in the Beaufort and Chukchi seas, Elementa: Science of the Anthropocene, 6, 48, https://doi.org/10.1525/elementa.304, 2018a.

355   Stern, H. L., Schweiger, A. J., Zhang, J., and Steele, M.: On reconciling disparate studies of the sea-ice floe size distribution, Elem Sci Anth, 6, 49, https://doi.org/10.1525/elementa.304, 2018b.

Wadhams, P., Squire, V. a., Goodman, D. J., Cowan, A. M., and Moore, S. C.: The attenuation rates of ocean waves in the marginal ice zone, Journal of Geophysical Research, 93, 6799, https://doi.org/10.1029/JC093iC06p06799, 1988.

Williams, T. D., Bennetts, L. G., Squire, V. A., Dumont, D., and Bertino, L.: Wave-ice interactions in the marginal ice zone.
360   Part 2: Numerical implementation and sensitivity studies along 1D transects of the ocean surface, Ocean Modelling, 71, 92–101, https://doi.org/10.1016/j.ocemod.2013.05.011, 2013.

Zhang, J., Schweiger, A., Steele, M., and Stern, H.: Sea ice floe size distribution in the marginal ice zone: Theory and numerical experiments, J. Geophys. Res. Oceans, 120, 3484–3498, https://doi.org/10.1002/2015JC010770, 2015.