Seasonal Sea Ice Presence Forecasting of Hudson Bay using Seq2Seq Learning

Nazanin Asadi [1] **Philippe Lamontage** [2] **Matthew King** [2 3] **Martin Richard** [2] **K Andrea Scott** [1]

Abstract

Accurate and timely forecasts of sea ice conditions are crucial for safe shipping operations in the Canadian Arctic and other ice-infested waters. Given the advancement of machine-learning methods and the recent observations on the declining trend of Arctic sea ice extent over the past decades due to global warming, new machine learning approaches are deployed to provide additional sea ice forecasting products. This study is novel in comparison with previous machine learning (ML) approaches in the sea-ice forecasting domain as it provides a daily spatial map of probability of ice presence in the domain up to 90 days. The predictions are further used to predict freeze-up/breakup dates and show their capability to capture both the variability and the increasing trend of open water season in the domain over the past decades.

1. Introduction

Sea ice presence is an important variable for northern communities in addition to offshore operations and shipping companies. Forecasting the presence of sea ice can be carried out at various spatial and temporal scales. Short-term forecasts at high spatial resolution are important for day-to-day operations and weather forecasting, whereas longer term (e.g. 60-90 day) forecasts are desired by shipping companies and offshore operators in the Arctic for strategic planning. Recent studies indicate a significant declining trend in Arctic sea ice extent over the past several decades in response to warming temperatures driven by climate change (Meier et al., 2014; Renner et al., 2014). This declining trend has resulted in increased shipping activities (Pizzolato et al., 2016) which requires accurate prediction of freeze-up/breakup dates and the open water season.

There is a wide body of literature on sea ice forecasting (Gue, 2016) with more recent approaches using machine learning (ML). Convolutional neural networks (CNNs) have been used to perform sea ice concentration prediction (Kim et al., 2020) with eight predictors composed of a sea ice concentration data and variables from reanalysis. The study trained 12 individual monthly models, and produced monthly spatial map of SIC. Their predictions were in good agreement with the minimum September ice extent from 2017. Similar to Kim et al. (2020), Horvath et al. (2020) focused on the September sea ice minimum, using a Bayesian logistic framework to predict both a monthly average sea ice concentration and an uncertainty. IceNet (Andersson et al., 2021) is a recent deep learning approach using a U-Net architecture with input from climate simulations and observational data to produce monthly sea ice concentration maps for the next 6 months. ML approaches have been compared with ensemble data assimilation at shorter time scales in Fritzner et al. (2020). However, none of the previously proposed ML approaches generate a forecast that propagates forward in time in a manner similar to a forecast model.

The goal of this study is to employ a sequence-to-sequence learning approach to provide a spatiotemporal forecast of the probability of sea ice at daily time scale over the region of Hudson Bay, with forecast lead times up to $\approx$ 90 days. The method is similar to operational forecasting studies (Chevallier et al., 2013) except i) we are using a data-driven statistical approach, as compared to a physics-based model; ii) our forecasted variable is the probability of ice at a grid location, as compared to sea ice concentration.

2. Data and Study Region

The present study utilizes ERA5 reanalysis for model predictors and validation. ERA5 is a recent reanalysis produced by the European Center (Hersbach & Dee, 2016). It consists of an atmospheric model coupled to the land, ice or ocean surface. The spatial resolution is $\approx$31 km with reanalysis fields available every hour from 1979 - present. Observations are assimilated into the atmospheric model using a

[*]Equal contribution [1]Department of Systems Design Engineering, University of Waterloo, Waterloo, Canada [2]Oceans, Coastal and River Engineering Research Centre, National Research Council Canada, Ottawa, Canada [3]Memorial University of Newfoundland, Newfoundland and Labrador, Canada. Correspondence to: Nazanin Asadi <n2asadi@uwterloo.ca>.

Tackling Climate Change with Machine Learning Workshop at ICML 2021.

4D-Variational data assimilation scheme. The current study utilizes the following input variables from ERA5 dataset over the period of 1985-2017: sea ice concentration, sea surface temperature, 2m temperature (T2M), surface sensible heat flux, wind 10 meter U-component (U10), wind 10 meter V-Component (V10), freezing degree days / melting degree days and landmask. All the input variables except sea ice concentration and landmask are normalized as the preprocessing process of the input data.

The ERA5 data are extracted over our study region, which corresponds to Hudson Bay, Hudson Strait and Foxe Basin. This region is mainly a seasonal ice zone, with ice melting completely each summer, except occasionally in Foxe Basin (northern portion of the domain). Hudson Bay and Hudson Strait are both home to several coastal communities. Hudson Strait is a region with year-round shipping to support natural resource extraction (Andrews et al., 2018).

3. Forecast model architecture

The medium-term forecasting problem of this study can be formulated as a spatiotemporal sequence forecasting problem that can be solved under the general sequence-to-sequence (Seq2Seq) learning framework in machine learning domain (Sutskever et al., 2014). In Seq2Seq learning, the target is to map a sequence of inputs to a sequence of outputs. The architecture of these models normally consist of 2 major components; encoder and decoder. The encoder component transforms a given input to an encoded state of fixed shape, while the decoder part takes that encoded state and generates an output sequence with the desired length. For this study, following this encoder-decoder architecture, a spatiotemporal sequence model is developed to predict a 90 day forecasts of ice presence probability given the last 3 days of historical inputs.

3.1. Basic model

The encoder section of the Basic network takes as input three rasters of daily environmental conditions. Each input sample is of size $(3 \times W \times H \times C)$ where 3 is the number of historical days, W and H are the width and height of the raster samples in their original resolution and C is the total number of input variables (here 8).

The encoder provides as output a single raster with the same height and width but higher number of channels such as to represent the fully encoded system state. First, each daily sample is passed through a feature pyramid network (Lin et al., 2017) so as to detect environmental patterns at both the local and large scales.

Next, the sequence of extracted feature grids are further processed through a convolutional LSTM layer (ConvLSTM) (Hochreiter & Schmidhuber, 1997; Xingjian et al., 2015), returning the last output state. This layer learns a single grid representation of the time series that also preserves spatial locality. Finally, the most recent day of historic input data is concatenated with the ConvLSTM output to produce a single data grid.

The final encoded state is then fed to a custom RNN decoder which extrapolates the state across the specified number of time-steps. It takes as input the encoded state with multiple channels and as output produces a state with the same width and height and desired time-steps.

The custom RNN decoder, as is common of many RNN layers, maintains both a cell state and a hidden state (Yu et al., 2019). First, the initial cell state and hidden state are initialized with the input encoded state. Then, at each time-step and for each of the states, the network predicts the difference, or residual, from the previous state to generate the updated states using depthwise separable convolutions (Howard et al., 2017). The output of the decoder section is the concatenation of the cell states from each time-step.

Finally, a time-distributed network-in-network (Lin et al., 2013) structure is employed to apply a 1×1 convolution on each time-step prediction to keep the grid size the same but reduce the number of channels to one, representing the daily probabilities of ice presence over the forecast period (e.g. up to 90 days).

3.2. Augmented model

A slight variant of the Basic model is developed so as to accept a second input. This second input corresponds to 3 variable climate normals over the required time-steps (e.g., 60 or 90), where these climate normal variables are the average of T2M, U10 and V10 from 1985 to the last training year for each forecast day. The original encoder structure for historical input data remains unchanged, but in addition, it is joined by a secondary feature pyramid network from climate normal. A secondary variant of the decoder component is implemented which accepts this encoded forecast sequence in order to produce superior estimates of the residuals at each of the future time-steps. Here, the decoder is designed in a way that the number of forecast input time-steps can be either equal or less than the forecast length. This network is referred to as the Augmented model.

4. Description of Experiments

For each month of a year a separate model is trained on data from the given month as well as the preceding and following month. For example, the 'April model' is trained using data from March 1 to May 31. This monthly model is initially trained on data from a fixed number of years, chosen to be 10 years. After this initial experiment, to predict each following test year i, we use a rolling forecast

prediction, meaning the model from year $i-1$ is retrained with data from year $i-2$ and validated on data from year $i-1$. For example, if the initial model is trained on 10 years, data from year 11 is used as validation and first predictions are launched at year 12. The model for year 12 is then retrained with data from year 11 and validated on year 12 to predict year 13 and so on. Thereby, the output's statistics are calculated on forecasts of 1996 to 2017. Since the retraining process only uses data of one year for training and validation, it is computationally fast and efficient. The models are implemented using the TensorFlow Keras open-source library with stochastic gradient descent (SD) optimizer with learning rate of 0.01, momentum of 0.9 and binary cross-entropy loss function.

In order to evaluate the performance of the neural network model, the results are compared with climate normal. This is defined as the average of ERA-5 sea ice concentration (thresholded at 15%) from 1985 to the last year in the training set for each experiment. While inputs of each model in training and test procedure are coming from 3 months of year, only the results from the central month (2nd of 3) is selected for evaluating the results of that model in the following section.

5. Results

5.1. Presence of Ice Forecasts

Given that our models predict spatial maps of sea ice presence over a grid at a spatial resolution of 31 km, we first apply a 50% threshold to this probability to convert each pixel to ice or water. From this, for each day in the test set, we have 90 binary accuracy maps of Hudson Bay. To summarize these results, in Fig 1(a,b,c) we show the model binary accuracy as a function of forecast lead day for each month. For example, the top row of Fig 1a shows the accuracy of forecasts launched in January for climate normal for forecast lead days of 1 to 90. E.g., the first top-left box in this figure (Fig 1(a)) corresponds to the accuracy after 1 day forecast for all forecasts launched between January 1 and January 31. The accuracy is very close to 100% for this month and for these lead times, as would be expected, because at this time the region is covered with ice. In contrast, for forecasts of June and July, the beginning of open water season, the climate normal struggles to accurately capture the ice cover for lead times of 1 to 50 days likely due to inter-annual variability and the impact of climate changes. However, the Basic and Augmented models proposed here have significantly higher accuracies than climate normal over these months (Fig 1d), especially in early lead times, with no degradation at longer lead times for Augmented model. We also note improvements in proposed models at early lead times for August, September, October and November, as compared to climate normal. When comparing the Augmented model with Climate normal (Fig 1e) and with the Basic model (Fig 1f), improvements in accuracy can be seen in particular for longer lead times in March/April and July/August, and improvements at shorter lead times (15-50 days) for November. Using climate normal forecasts as additional input for Augmented model is showing its impact here where in the periods that Basic model is worse than climate normal (Fig 1d), the Augmented model has better and closer accuracy to climate normal (Fig 1e, Fig 1f).

5.2. Assessment of operational capability

The accuracy of the model predictions in predicting freeze-up/breakup date and their comparison with climate normal is an indicative of operational capability of the trained models. The freeze-up date of each pixel in a year is defined as the first date in the freeze-up season (Oct 1st to Jan 31st for Hudson Bay) that ice (value of 1) is observed for 15 continuous days. A similar procedure is carried out to define breakup, with the exception that the pixel must be considered water (value of 0) for 15 continuous days in the breakup season (May 1st to July 31st for Hudson Bay) for breakup to have occurred. Accordingly, open water season of each year is the number of days between its breakup and freeze-up date.

Figure 2 represents the capability of the models at 30 and 60 lead day in capturing the changes in open water season duration versus climate normal. Each map is showing the changes (number of days) between the median of open water season duration of second decade (2007-2017) and first decade (1996-2006) in the test results. In the northern part of the domain in Foxe Basin and Gulf of Boothia, the open-water season is too short or inexistent to be captured by the methodology. However, the figure shows that there is an increase of up to 30 days in the open water season of the southern and western sections of Hudson Bay. While climate normal is not able to show this increase, the models show more promising forecasts, especially at 30 lead day.

6. Conclusion

The purpose of this study was to propose a ML method to provide daily forecast maps of sea ice presence probability up to 90 days in advance. The binary accuracy of the proposed models represented improvements of up to 10% relative to climate normal for breakup and freeze-up season, especially for early lead days. The analysis on increasing trend of open water season over two decades indicated the ability of the proposed models at 30 lead day to capture the trend in southern and western part of the region while the climate normal was not able to show such trend. As future work, we plan to compare this approach to another developed for this region using canonical correlation analysis (Tivy et al., 2011), expand the experiments over the entire Arctic region, and deploy ensemble methods and more

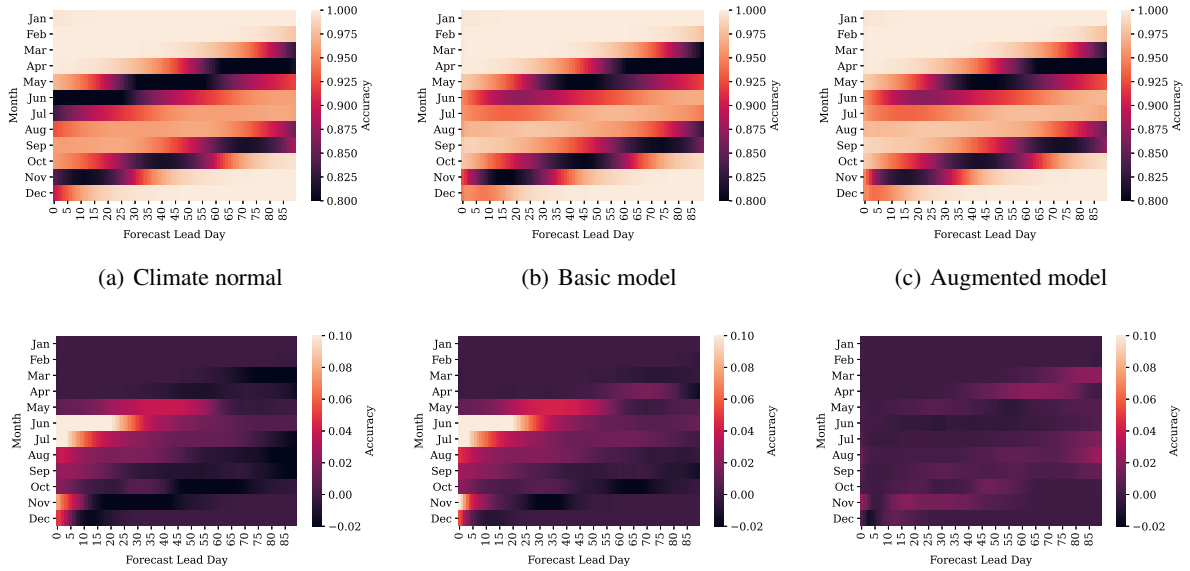

(a) Climate normal (b) Basic model (c) Augmented model

(d) Basic model-Climate normal (e) Augmented model-Climate normal (f) Augmented model-Basic model

Figure 1. Model performance and improvements as a function of lead time. Top row panels describe performance of each model (a)-(c) while bottom row represents the accuracy differences between the models (d-f). Most differences are observed in breakup and freeze-up seasons.

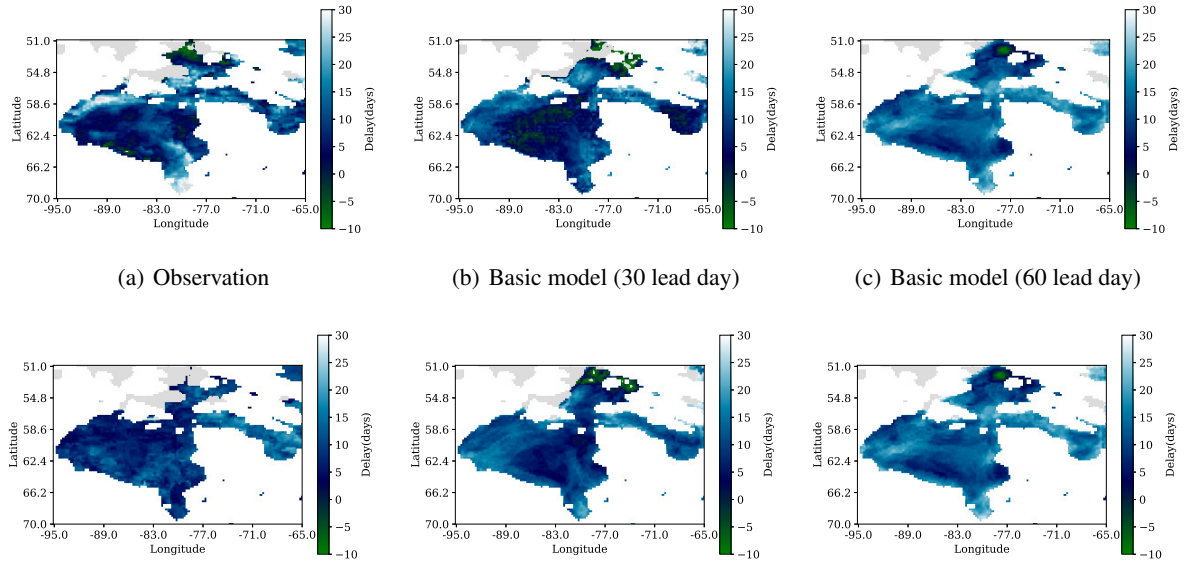

(a) Observation (b) Basic model (30 lead day) (c) Basic model (60 lead day)

(d) Climate normal (e) Augmented model (30 lead day) (f) Augmented model (60 lead day)

Figure 2. Difference between the median of the open water season length of two decades (1996-2006 vs 2007-2017) in terms of number of days for Hudson Bay. The grey zone in the north is the area where the open water season is too short or inexistent to be captured by the methodology. Despite climate normal, the proposed models, especially at 30 lead day, are able to capture the trend in southern and western part of the region.

recent deep learning architectures.

7. Acknowledgements

The authors would like to acknowledge funding from the National Research Council of Canada through the Ocean and AI4Logistics Programs and computing resources provided by Compute Canada.

References

A review on Arctic sea-ice predictability and prediction on seasonal to decadal time scales. *Quarterly Journal of the Royal Meteorological Society*, 142:546–561, 2016.

Andersson, T. R., Hosking, J. S., Pérez-Ortiz, M., Paige, B., Elliott, A., Russell, C., Law, S., Jones, D. C., Wilkinson, J., Phillips, T., et al. Seasonal Arctic sea ice forecasting with probabilistic deep learning. 2021.

Andrews, J., Babb, D., Barber, D. G., and Ackley, S. F. Climate change and sea ice: shipping in Hudson Bay, Hudson Strait, and Foxe Basin (1980–2016). *Elementa: Science of the Anthropocene*, 6, 2018.

Chevallier, M., Salas y Mélia, D., Voldoire, A., and Deque, M. Seasonal forecasts of pan-Arctic sea ice extent using a GCM-based seasonal prediction system. *Journal of Climate*, 26:6092–6104, 2013.

Fritzner, S., R., G., and Christensen, K. Assessment of high resolution dynamical and machine learning models for prediction of sea ice concentration in a regional application. *Journal of Geophysical Research, Oceans*, 2020.

Hersbach, H. and Dee, D. ERA5 reanalysis is in production. *ECMWF newsletter*, 147(7):5–6, 2016.

Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

Horvath, S., Stroeve, J., Rajagopalan, B., and Kleiber, W. A bayesian logistic regression for probabilistic forecasts of the minimum september arctic sea ice cover. *Earth and Space Science*, 7(10):e2020EA001176, 2020.

Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., and Adam, H. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.

Kim, Y., Kim, H.-C., Han, D., Lee, S., and Im, J. Prediction of monthly Arctic sea ice concentrations using satellite and reanalysis data based on convolutional neural networks. *The Cryosphere*, 14:1083–1104, 2020.

Lin, M., Chen, Q., and Yan, S. Network in network. *arXiv preprint arXiv:1312.4400*, 2013.

Lin, T.-Y., Dollár, P., Girshick, R., He, K., Hariharan, B., and Belongie, S. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2117–2125, 2017.

Meier, W. N., Hovelsrud, G. K., Van Oort, B. E., Key, J. R., Kovacs, K. M., Michel, C., Haas, C., Granskog, M. A., Gerland, S., Perovich, D. K., et al. Arctic sea ice in transformation: A review of recent observed changes and impacts on biology and human activity. *Reviews of Geophysics*, 52(3):185–217, 2014.

Pizzolato, L., Howell, S. E., Dawson, J., Laliberté, F., and Copland, L. The influence of declining sea ice on shipping activity in the canadian arctic. *Geophysical Research Letters*, 43(23):12–146, 2016.

Renner, A. H., Gerland, S., Haas, C., Spreen, G., Beckers, J. F., Hansen, E., Nicolaus, M., and Goodwin, H. Evidence of arctic sea ice thinning from direct observations. *Geophysical Research Letters*, 41(14):5029–5036, 2014.

Sutskever, I., Vinyals, O., and Le, Q. V. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pp. 3104–3112, 2014.

Tivy, A., Howell, S., Alt, B., Yackel, J., and Carrieres, T. Origins and levels of seasonal skill for sea ice in Hudson Bay using canonical correlation analysis. *Journal of Climate*, 24:1378–1394, 2011.

Xingjian, S., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-K., and Woo, W.-c. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*, pp. 802–810, 2015.

Yu, Y., Si, X., Hu, C., and Zhang, J. A review of recurrent neural networks: Lstm cells and network architectures. *Neural computation*, 31(7):1235–1270, 2019.