

THESIS

AI-INFORMED MODEL ANALOGS FOR SUBSEASONAL-TO-SEASONAL PREDICTION

Submitted by

Jacob B. Landsberg

Department of Atmospheric Science

In partial fulfillment of the requirements

For the Degree of Master of Science

Colorado State University

Fort Collins, Colorado

Summer 2025

Master's Committee:

Advisor: Elizabeth Barnes

Russ Schumacher

Jay Ham

Copyright by Jacob Landsberg 2025

All Rights Reserved

ABSTRACT

AI-INFORMED MODEL ANALOGS FOR SUBSEASONAL-TO-SEASONAL PREDICTION

Subseasonal-to-seasonal forecasting is crucial for public health, disaster preparedness, and agriculture, and yet it remains a particularly challenging timescale to predict. We explore the use of an interpretable AI-informed model analog forecasting approach, previously employed on longer timescales, to improve S2S predictions. Using an artificial neural network, we learn a mask of weights to optimize analog selection and showcase its versatility across three varied prediction tasks: 1) classification of Week 3-4 Southern California summer temperatures; 2) regional regression of Month 1 midwestern U.S. summer temperatures; and 3) classification of Month 1-2 North Atlantic wintertime upper atmospheric winds. The AI-informed analogs outperform traditional analog forecasting approaches, as well as climatology and persistence baselines, for deterministic and probabilistic skill metrics on both climate model and reanalysis data. We find the analog ensembles built using the AI-informed approach also produce better predictions of temperature extremes and exhibit more reliable forecast uncertainty. Finally, by using an interpretable-AI framework, we analyze the learned masks of weights to better understand S2S sources of predictability.

ACKNOWLEDGEMENTS

I would like to thank my advisor Libby Barnes for her guidance and support throughout this project, as well as Matt Newman and Martin Fernandez for their helpful feedback and insightful comments. I would also like to thank Russ Schumacher and Jay Ham for serving on my committee. I also acknowledge support from NOAA grants #NA22OAR4310621 and #NA19OAR4590151. Finally, I would like to thank my family, friends, and group members for their encouragement and support.

TABLE OF CONTENTS

	ABSTRACT	ii
	ACKNOWLEDGEMENTS	iii
	LIST OF TABLES	v
	LIST OF FIGURES	vi
Chapter 1	Introduction	1
Chapter 2	AI-Informed Model Analogs for Subseasonal-to-Seasonal Prediction	3
2.1	Introduction to Analog Forecasting	3
2.2	Methods	5
2.2.1	Prediction Tasks	5
2.2.2	AI-Informed Analog Approach	5
2.2.3	Artificial Neural Network to Learn Mask of Weights	7
2.2.4	Data	7
2.2.5	Metrics	9
2.3	Results	10
2.3.1	Week 3-4 Windows of Opportunity in Southern California	10
2.3.2	Month 1 Temperature Extremes Over the Midwestern U.S.	13
2.3.3	Month 1-2 Mask Exploration in the North Atlantic	15
2.4	Conclusions	17
Chapter 3	Discussion and Future Work	22
Appendix A	Data and Code Availability	31
Appendix B	Neural Network Details	32
B.1	Hyperparameters	32
B.2	Constrained Inverse L_2 Regularization	32
Appendix C	Data: CESM2-LE Members	34
Appendix D	Figures	35

LIST OF TABLES

2.1	The three prediction tasks.	5
B.1	Hyperparameters used in the model training.	32
C.1	List of monthly ensemble members used for the analog library, training SOIs, validation SOIs, and testing SOIs.	34
C.2	List of daily ensemble members used for the analog library, training SOIs, validation SOIs, and testing SOIs.	34

LIST OF FIGURES

2.1	Schematic of the steps in the AI-informed analog approach: 1) Assemble the SOI and the library of d potential analogs. 2) Multiply all potential analogs and the SOI by the learned mask of weights. 3) Compute the MSE between the weighted SOI and the d weighted potential analogs, and select the k closest analogs (in this example $k = 2$). 4) Find the values of the analogs in the target region after the desired lead time. 5) Use the target field of the k closest analogs' evolution in the target region as the prediction for the SOI (this example is a regression problem, so the mean value is taken).	6
2.2	Schematic of the neural network setup to learn the weighted mask. One SOI and one analog are multiplied by a layer of learnable weights. The MSE between the two weighted inputs is computed and passed through a linear scaling layer. This output represents the predicted difference in the two maps' targets. Loss is computed as the MSE between the predicted difference of the targets and the true difference of the targets.	8
2.3	The learned mask for Task #1, Southern California summer temperature classification. The cyan box outlines the target region.	11
2.4	Skill scores for a) CESM2-LE accuracy, b) CESM2-LE BS, c) ERA5 accuracy, and d) ERA5 BS for Week 3-4 Southern California temperature classification.	12
2.5	Discard plot based on ensemble agreement for Week 3-4 Southern California temperature classification using 1500 analogs, testing on ERA5 data.	13
2.6	The learned mask for Task #2, midwestern U.S. summer temperature regression. The cyan box outlines the target region.	14
2.7	Skill scores for a) CESM2-LE accuracy, b) CESM2-LE CRPS, c) ERA5 accuracy, and d) ERA5 CRPS for Month 1 midwestern U.S. temperature regression.	15
2.8	Discard plot based on predicted extremity with an ensemble of 50 analogs for midwestern U.S. summer temperature regression, testing on ERA5 data.	16
2.9	The learned mask for Task #3, North Atlantic winter U250 classification. The cyan box outlines the target region.	19
2.10	Skill scores for a) CESM2-LE accuracy, b) CESM2-LE BS, c) ERA5 accuracy, and d) ERA5 BS for Month 1-2 North Atlantic U250 classification.	20
2.11	Examples of masks with each of the ablation methods: (a) Ablating entire fields, (b) Threshold ablation, and (c) Ablating specific regions.	21
2.12	BS skill for different ablation methods evaluated on CESM2-LE data.	21
D.1	Discard plot for Week 3-4 Southern California temperature classification using 1500 analogs, testing on CESM2-LE data. Unlike when testing on reanalysis data, global ensemble agreement serves as a reasonable metric for forecast uncertainty. However, the learned-mask ensemble agreement still showcases a greater decrease in error for states with higher ensemble agreement.	35
D.2	With regression problems, as the number of analogs approaches the total library size, the prediction becomes more and more similar to a climatological prediction. In this example the library size is ~ 10000 , where the skill is 0.	36

D.3	A discard plot for an ensemble of 50 analogs for midwestern U.S. summer temperature regression, testing on CESM2-LE data. While the regional mask perform relatively better on extreme predictions for CESM2-LE compared to ERA5 data, there is still a much larger improvement for extreme predictions with the learned mask.	36
D.4	In Task #3 a field of values rather than the average value across the region of interest is predicted. This map shows the average skill across SOIs when using a 400 analog ensemble on ERA5 data. At almost every grid point, the learned mask outperforms the climatological prediction.	37
D.5	Using inverse L_2 regularization with $\lambda_2 = 100$ for the North Atlantic (Task #3) results in a sparser map.	38
D.6	Expanded table, showing changes in BS skill with different ablation methods, for 100-800 analogs	39

Chapter 1

Introduction

Producing more accurate weather and climate forecasts can save lives and money. For example, the World Bank estimates that if all low- and middle-income countries had hydrometeorological information comparable to that of Europe, enabling better forecasts, over 20,000 lives and between \$3.3 and \$32 billion could be saved annually [17]. Similarly, as climate change intensifies, improved climate predictions can help policymakers allocate resources more effectively for adaptation and mitigation efforts [13]. However, while forecasters have made significant strides improving predictions on both shorter-term weather (<2 weeks) [4] and longer-term climate (>2 months) [3] timescales, critical gaps remain in our ability to forecast at timescales in between weather and climate regimes [34, 6, 39, 44].

Forecasting between 2 weeks and ~ 2 months, often referred to as Subseasonal-to-Seasonal prediction (S2S), is essential for public health, disaster preparedness, agriculture, and energy and water management [44]. Despite the clear benefits of skillful predictions at these timescales, S2S forecasting remains particularly challenging. Often described as a ‘predictability desert’ [34, 6], S2S forecasts cannot rely solely on initial atmospheric conditions, as is common in short-term numerical weather prediction, nor on the slow-varying boundary conditions that underpin climate outlooks [34, 39]. Instead, forecasters must integrate information from initial conditions, boundary conditions, and S2S modes of variability, such as the Madden-Julian Oscillation (MJO) [47], to improve prediction skill [39]. However, on S2S timescales, the strength of these predictability sources and their teleconnections (i.e., long-distance relationships between different climate phenomena) remain uncertain [27, 39], and forecast skill—such as the accuracy of summertime surface temperature predictions in North America—remains relatively low [5, 31]. A variety of tools have been used to approach the S2S forecasting challenge. Dynamical models have slowly but steadily improved S2S forecast skill [32] and data-driven approaches, like fully-AI models, can now forecast phenomena such as the North Atlantic Oscillation (NAO) and MJO at S2S lead times

[20, 6] with similar skill to dynamical models. To further improve forecasts, there has recently been a renewed focus on pinpointing climate states that represent times of enhanced predictability [24, 25, 2]. Identifying these ‘windows of opportunity’ is a potential approach to improve skill on S2S timescales by allowing forecasters to know when forecast uncertainty is high or when they can leverage these times of enhanced predictability for more accurate forecasts [24].

Chapter 2

AI-Informed Model Analogs for Subseasonal-to-Seasonal Prediction

2.1 Introduction to Analog Forecasting

Here, we tackle S2S prediction by combining a variety of these existing methodologies and employing an AI-informed model analog forecasting approach. Analog forecasting rests on the premise that climate states with similar initial conditions tend to evolve in a consistent manner [21, 48]. By identifying past states resembling current conditions, their subsequent evolution can offer plausible trajectories for future conditions. For a variety of forecasts, from the tropics to the northern high latitudes, analog forecasting has been shown to rival the skill of global climate models [22, 8, 41] all while offering several key advantages. Unlike fully-AI models, analog forecasting is intuitive, interpretable, and can uphold physical laws [33, 9]; moreover, compared to global dynamic climate models, analog forecasting is highly computationally efficient [8].

Analogs offer an interpretable, physical model that is helpful for diagnosing errors and probing physical drivers, while their fast computational speed allows for the quick generation of ensembles of forecasts. Creating proficient ensembles is a key way to improve skill on S2S timescales [14, 29, 16], provide probabilistic forecasts [28, 19, 43], and even help explore windows of opportunity [19, 43]—essential on S2S timescales. For instance, with a calibrated ensemble of forecasts, one can use ensemble member agreement as a sign of a lower forecast uncertainty to identify windows of opportunity [12]. However, despite these advantages in computation and interpretability, successful analog forecasting hinges on having both a robust library of analogs and a reliable method to identify sufficiently similar past states.

To address this need for a large analog library, we turn to climate models, which have orders of magnitude more climate realizations than we have observational data [9, 26]. Yet, even with

climate models, finding perfect analogs is impractical—estimates suggest over 10^{30} years of data would be needed to match two atmospheric flow stream patterns in just the Northern Hemisphere within observational error [38]. Hence, determining the conditions that make a climate state an adequately close analog, rather than a perfect one, is crucial. For example, regional matching has been used to identify close analogs for seasonal tropical Indo-Pacific Ocean prediction [9]; global matching has been used for multi-decadal global predictions [23]; and area-specific matching has been used for annual-to-multi-year Pacific Decadal Oscillation prediction [46]. These methods for selecting analogs have been shown to work for certain problems, although they demand either a huge library of analogs (as in global matching) or depend on prior knowledge of physical drivers and teleconnections (as in regional or area-specific matching).

Here, we explore an alternative, AI-based spatial weighting approach originally introduced by [33]. We train a neural network to output a mask of weights that highlights where it is most important for initial conditions to match, such that two states will evolve similarly. Using a learned set of weights to find optimal analogs reduces reliance on prior knowledge and enables investigation of which regions and variables are most essential for two climate states to follow similar future trajectories. This method of optimized analog forecasting was first successfully applied to annual-to-decadal sea surface temperature prediction [33] and has since been extended to multi-year-to-decadal 2-meter temperatures [11] and El Niño-Southern Oscillation (ENSO) predictions [37]. Importantly, while [37] showed skill with this AI-informed analog approach for ENSO predictions on seasonal-to-interannual timescales, they were unable to achieve skill on S2S timescales.

Here, we show that this AI-based analog forecasting approach can achieve skill beyond traditional analog methods on S2S timescales while maintaining interpretability and computational efficiency. We highlight the benefits of using AI-based analogs across three varied prediction tasks: 1) classification of Week-3-4 Southern California summer temperatures; 2) regional regression of Month-1 midwestern U.S. summer temperatures; and 3) classification of Month-1-to-2 North Atlantic wintertime upper atmosphere winds. Through these three prediction tasks we show the AI-based analog approach outperforms traditional analog forecasting approaches, as well as cli-

matology and persistence baselines on reanalysis data on S2S timescales. Further, we exploit analog ensembles to quantify forecast uncertainty, and, by leveraging an interpretable AI-forecasting framework, analyze the learned masks of weights to better understand S2S sources of predictability.

2.2 Methods

2.2.1 Prediction Tasks

We demonstrate the skill of our analog forecasting approach for three prediction tasks described in Table 2.1, opting for a varied set of examples to test the generalizability of the method across different S2S prediction problems. We apply the AI-informed analog approach to both classification and regression tasks, to different regions, seasons, variables, and lead times. Each of these prediction tasks has a unique learned mask of weights to optimize the choice of analogs.

Table 2.1: The three prediction tasks.

Prediction Task	1	2	3
Region	Southern California	Midwestern U.S.	North Atlantic
Data Frequency	Smoothed Daily	Monthly	Monthly
Prediction Time	Week 3-4	Month 1	Month 1-2
Prediction Season	June Week 3 - Sept. Week 3	July - Sept.	Dec. - Feb.
Input \rightarrow Target	Skin Temp \rightarrow Skin Temp	Skin Temp \rightarrow Skin Temp	Skin Temp + U250 \rightarrow U250
Target Type	Single Value	Single Value	Field
Classification/Value	Tercile Classification	Regression	Tercile Classification

2.2.2 AI-Informed Analog Approach

Most traditional analog forecasting methods follow a similar approach: to predict how a certain climate state (referred to as a state of interest or SOI) will evolve, one finds the closest k matches (for $k \geq 1$, in the analog) library. ‘‘Closeness’’ is often measured by minimizing a distance measure,

like mean-squared error (MSE), between the SOI and potential analogs either across the entire globe, or across a region of interest. One can then use the trajectories of these closest matches as a prediction for how the SOI will evolve into the future. Rather than predicting the evolution of the whole globe, one often evaluates the analog skill in a specific region of interest, which we refer to as the target region.

We take a similar approach, except here we utilize a soft mask of weights to measure closeness between the SOI and potential analogs. Prior to computing the distance measure between the SOI and each potential analog, we multiply the entire library and the SOI by a learned mask of weights (Steps 1 and 2 in Figure 2.1). This mask, therefore, highlights or dampens the importance of conditions matching in certain areas of the globe for a potential analog to be considered close to the SOI. We then use MSE to compute the closest analogs after weighting, selecting the k closest analogs (Step 3 in Figure 2.1). Lastly, we use the k closest analogs' mean evolution (for regression problems) or majority vote (for classification) in the target region as our final prediction (Steps 4 and 5 in Figure 2.1).




Figure 2.1: Schematic of the steps in the AI-informed analog approach: 1) Assemble the SOI and the library of d potential analogs. 2) Multiply all potential analogs and the SOI by the learned mask of weights. 3) Compute the MSE between the weighted SOI and the d weighted potential analogs, and select the k closest analogs (in this example $k = 2$). 4) Find the values of the analogs in the target region after the desired lead time. 5) Use the target field of the k closest analogs' evolution in the target region as the prediction for the SOI (this example is a regression problem, so the mean value is taken).

2.2.3 Artificial Neural Network to Learn Mask of Weights

The weighted mask is learned by an artificial neural network that is similar to that of [33], with minor modifications to its final layers. The network, depicted in Figure 2.2, takes two maps as input. These maps are both multiplied by a grid of learnable weights of the same size as the inputs, resulting in two weighted maps. The mask is restricted to have a mean of 1 across all weights, such that during training the mask effectively moves weight between different areas of the globe. The MSE between these two weighted maps is passed through a single linear scaling layer. The output of this layer represents the network’s prediction of the MSE between the two maps in the target region after they have evolved (i.e., after the desired lead time). Loss is computed as the MSE between the predicted difference of the targets and the true difference of the targets. Details of the network setup and hyperparameters can be found in Appendix B.

Our network deviates slightly from that of [33], in that we use a single linear layer instead of multiple dense layers at the end of the network. We restrict the linear layer’s weight to be ≥ 0 to ensure a monotonically increasing relationship between the MSE of the maps’ weighted inputs and the predicted MSE of the maps’ target regions. This better matches our process for selecting analogs as described in Section 2.2.2, as we expect that two maps with a smaller weighted MSE will also evolve to have a smaller MSE between their targets. This switch to a single linear layer resulted in a negligible change in skill, but increased network parsimony and training speed.

2.2.4 Data

We use output from the Community Earth System Model 2–Large Ensemble (CESM2-LE) [7] in order to have a sufficiently large analog library and to learn the weighted mask. As we will show, the AI-informed analog approach produces skillful predictions when evaluated on both CESM2-LE data, in a perfect-model framework, and on ECMWF Reanalysis v5 (ERA5) data [15]. We employ a 7-day sliding window to smooth daily CESM2-LE and ERA5 data, using a backward moving average for input data and a forward moving average for target data, and also make use of monthly-mean fields. All smoothed daily data is regridded via bilinear interpolation to $2.5^\circ \times$

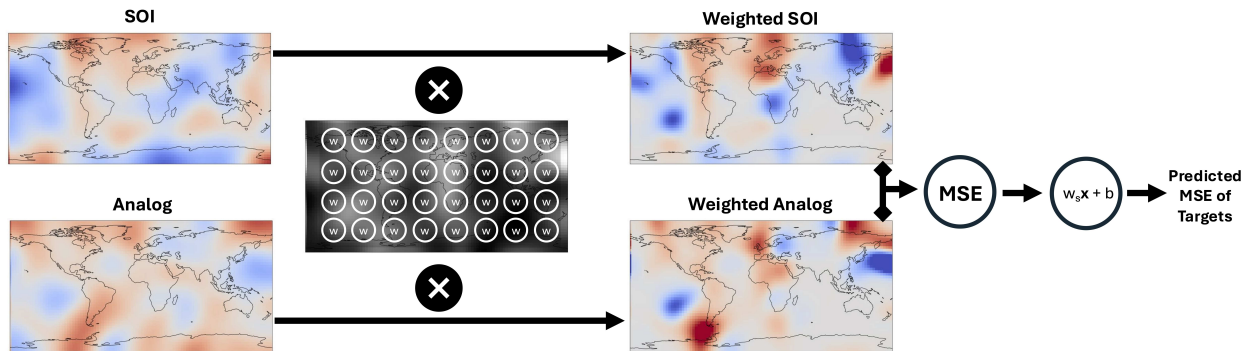


Figure 2.2: Schematic of the neural network setup to learn the weighted mask. One SOI and one analog are multiplied by a layer of learnable weights. The MSE between the two weighted inputs is computed and passed through a linear scaling layer. This output represents the predicted difference in the two maps’ targets. Loss is computed as the MSE between the predicted difference of the targets and the true difference of the targets.

2.5° resolution, while monthly data is resolved at .25° x .24° (natively for CESM2-LE data, and bilinearly interpolated for ERA5 data). As our analog library of daily data climate maps is $\sim 10\times$ larger than the library of monthly data, we use a coarser resolution for the daily data to reduce the memory load. For all data sources, we convert the data to anomalies about the climatological seasonal cycle and then to standard deviations at each grid point. However, between data sources, we handle the anthropogenic effects of climate change slightly differently, as will be discussed next.

CESM2-LE Data

We use monthly CESM2-LE data from 1850-2100 that employs CMIP6 historical and SSP3-7.0 future radiative forcing scenarios [35]. We take all 100 members to calculate the ensemble mean, which we subtract from each individual member to both remove the effects of anthropogenic climate change and to convert the data to anomalies from the seasonal cycle. To increase speed and reduce memory load, we then use only a third of the members for training and the analog library. These members are divided between the analog library and SOIs, with fields from 19 members

composing the library and fields from 14 members serving as the SOIs (see Table C.1 for member details). We partition the SOIs with a 10/2/2 member split for training, validation, and testing respectively.

For the daily CESM2-LE data, we subtract the linear trend from each calendar day at each grid point. We include a shorter timespan of data from each member (1850-1949) than the monthly data, as each daily-data year contains more than $30\times$ the amount of samples. The analog library is composed of fields from the first 5 members, while fields from the next 4 members (with a 2/1/1 training/validation/testing split) make up the SOIs (see Table C.2 for member details).

ERA5 Data

We use ERA5 daily data from 1942-2023 and monthly data from January 1940 to July 2024. For both daily and monthly ERA5 data, we fit and subtract a third-order polynomial at each grid point and each calendar day, or calendar month, respectively to define detrended anomalies from the seasonal cycle. The ERA5 data acts as a second test set to evaluate skill on observations.

2.2.5 Metrics

We employ deterministic and probabilistic error metrics for each type of prediction task (i.e., classification and regression). For regression (Task #2) we compute mean absolute error (MAE) and continuous ranked probability score (CRPS).

MAE is defined as

$$MAE = \frac{1}{N} \sum_{i=1}^N |f_i - o_i| \quad (2.1)$$

where N is the number of samples, f_i is the predicted value for sample i , and o_i is the true value for sample i .

CRPS is defined as

$$CRPS(F, x) = \int_{-\infty}^{\infty} (F(y) - H(y - x))^2 dy \quad (2.2)$$

where $F(y)$ is the cumulative distribution function of the forecast, x is the true value, and $H(y-x)$ is the Heaviside step function, which is 0 for $y < x$ and 1 for $y \geq x$. CRPS ranges from 0 (for a perfect forecast) to ∞ .

For classification (Tasks #1 and #3), we compute misclassification rate and the multi-class Brier Score (BS).

Misclassification rate is defined as

$$\text{Error Rate} = \frac{\text{Number of Incorrect Classifications}}{\text{Total Number of Predictions}} \quad (2.3)$$

BS is defined as

$$BS = \frac{1}{N} \sum_{i=1}^N \sum_{k=1}^K (f_{ik} - o_{ik})^2 \quad (2.4)$$

where N is the number of samples, K is the number of classes, f_{ik} is the predicted probability for class k for sample i , and o_{ik} is the true value (1 if the true class is k , otherwise 0). BS ranges from 0 (for a perfect forecast) to 2.

We convert all types of error to skill scores by comparing them to the error of a climatological forecast:

$$\text{Skill Score} = 1 - \frac{\text{Error}}{\text{Error}_{\text{climatology}}} \quad (2.5)$$

All skill scores are strictly ≤ 1 , with a skill score of 1 indicating perfect skill and a skill score of 0 indicating equal skill to a climatological forecast. A negative skill score indicates worse skill than a climatological forecast.

2.3 Results

2.3.1 Week 3-4 Windows of Opportunity in Southern California

We first assess the short-range S2S skill of the AI-based analog approach by classifying Week 3-4 Southern California ($32^\circ - 37^\circ\text{N}$, $116^\circ - 121^\circ\text{W}$) summer temperatures (Task #1). The three target classes (cold, neutral, and warm) are formed by splitting the target temperatures into terciles,

ensuring all classes are equally sized. Terciles for classifying the analog library are determined using the data within the analog library and terciles for the test set are defined based on the data in the test set to limit the impact of CESM2-LE biases relative to ERA5. We predict each 2-week period from the third week of June through the third week of September, using the learned weights in Figure 2.3. This mask exhibits weights that are distributed globally, yet unevenly, with noticeably increased weight around the western U.S. as well as in the North Pacific.

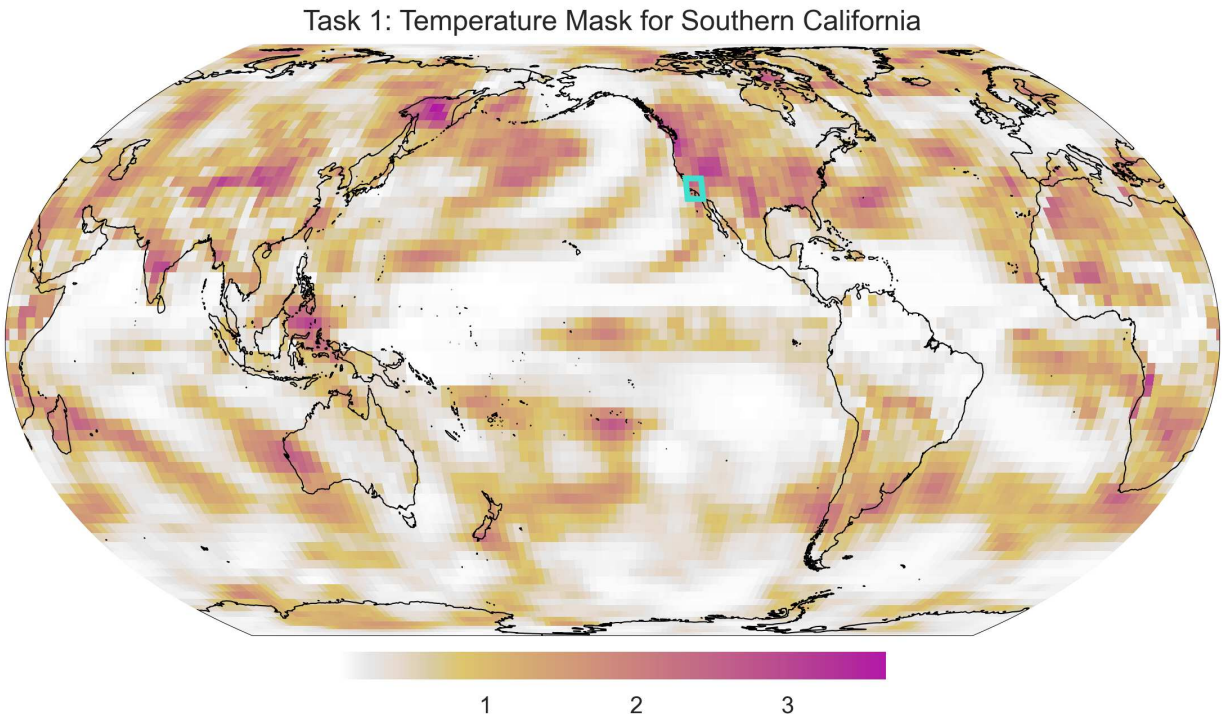


Figure 2.3: The learned mask for Task #1, Southern California summer temperature classification. The cyan box outlines the target region.

We include a regional and a global analog baseline in addition to persistence and climatological baselines to evaluate the relative skill of the learned mask approach. To create a global baseline, we select analogs by matching conditions over the entire globe (equivalent to a weighted mask of 1s everywhere). We create a regional baseline by selecting analogs via matching conditions only in the target region (equivalent to a weighted mask of 1s in the target region and 0s everywhere else). With the learned mask, MAE and BS skill scores exceed other baselines when testing both

on CESM2-LE and ERA5 data (Figure 2.4). With the CESM2-LE test set, the skill score peaks at 2000 analogs, while for ERA5, the skill score peaks at 1500 analogs.

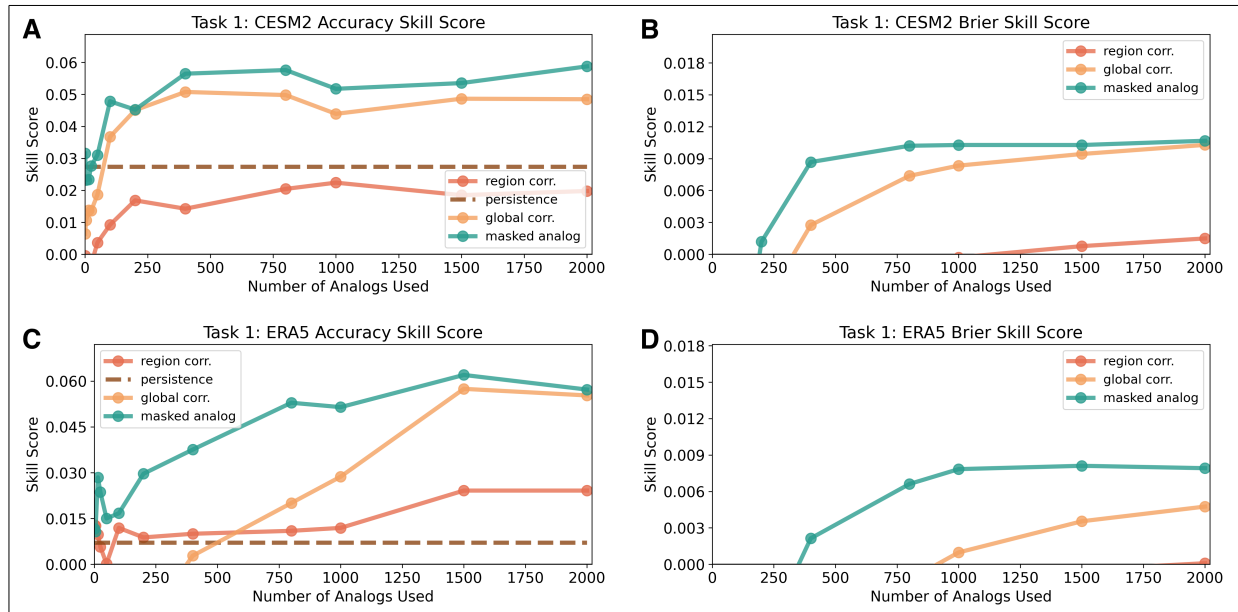


Figure 2.4: Skill scores for a) CESM2-LE accuracy, b) CESM2-LE BS, c) ERA5 accuracy, and d) ERA5 BS for Week 3-4 Southern California temperature classification.

We diagnose whether the analog ensembles can offer insights into windows of opportunity via discard plots. Figure 2.5 uses ERA5 data and a 1500-analog ensemble to show the change in misclassification rate from a climatological forecast as samples with lower ensemble agreement are discarded (for CESM2 data see Figure D.1). Here, ensemble agreement is computed as the fraction of ensemble members that agree on the majority prediction. We see that over all samples the mask offers just over a 4% reduction in misclassification rate relative to climatology, but this reduction grows essentially monotonically to over 9% for the $\sim 25\%$ of samples with the highest ensemble agreement. This is not the case with a global mask's ensemble, which does not exhibit as precipitous of a reduction in misclassification rate and actually *increases* in misclassification rate until the $\sim 50\%$ cutoff mark.




Figure 2.5: Discard plot based on ensemble agreement for Week 3-4 Southern California temperature classification using 1500 analogs, testing on ERA5 data.

2.3.2 Month 1 Temperature Extremes Over the Midwestern U.S.

We also explore how well the AI-informed analog forecasting approach can perform regression, by predicting monthly summer midwestern U.S. ($36^{\circ} - 49^{\circ}\text{N}$, $90^{\circ} - 106^{\circ}\text{W}$) temperatures (Task #2) with a focus on extremes. We predict the temperature (in standard deviations— σ) each month from July through September, using the learned mask in Figure 2.6. The mask displays a strong emphasis on the target region and preferential weighting in the mid-latitudes of the Northern Hemisphere as well as the Maritime Continent. Moreover, compared to the mask in Figure 2.3, the mask in Figure 2.6 has a more concentrated distribution of weights, highlighting the regional importance of the central U.S. for predicting midwestern summer temperatures.

As in Task #1, the learned mask outperforms all baselines (Figure 2.7). All skill scores peak at 50 analogs for both CESM2-LE and ERA5 data, except for CESM2-LE CRPS, which peaks at 100 analogs. The number of analogs in the ensembles are much lower than in Task #1, since we have moved from a 3-class classification problem to a regression problem. With a regression

Task 2: Temperature Mask for the Midwestern U.S.

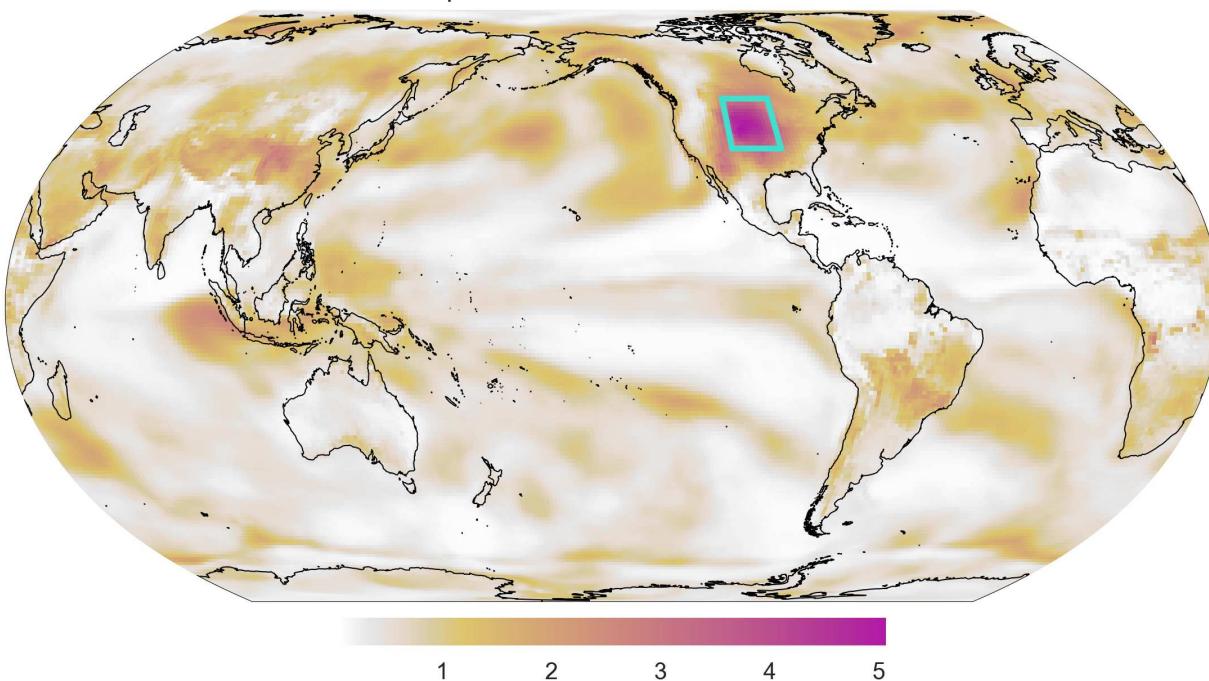


Figure 2.6: The learned mask for Task #2, midwestern U.S. summer temperature regression. The cyan box outlines the target region.

problem, if the number of analogs in the ensemble is too high, the ensemble mean will converge to climatology. An example of this regression to the mean can be found in Figure D.2.

While overall improvement in temperature forecasting on S2S timescales is important, better prediction of extreme temperatures has an outsized impact on enhancing agricultural production, public health, and energy management [10]. Thus, we focus on assessing the AI-based analog's ability to predict extreme temperatures. We again utilize discard plots, where we compare how MAE changes relative to a climatological prediction for more extreme samples. We denote extremity simply as the absolute value of the prediction, i.e. a measure of how far from climatology the prediction is. In Figure 2.8, we show the discard plot with ERA5 data using an ensemble of 50 analogs. The AI-based analogs exhibits a marked decrease in relative error for samples with more extreme predictions. As we show error relative to climatology, it may be unsurprising that the AI-informed analogs would have lower relative error on more extreme events. However, this is not the case for the regional baseline, where there is only a slight decrease in error for the most

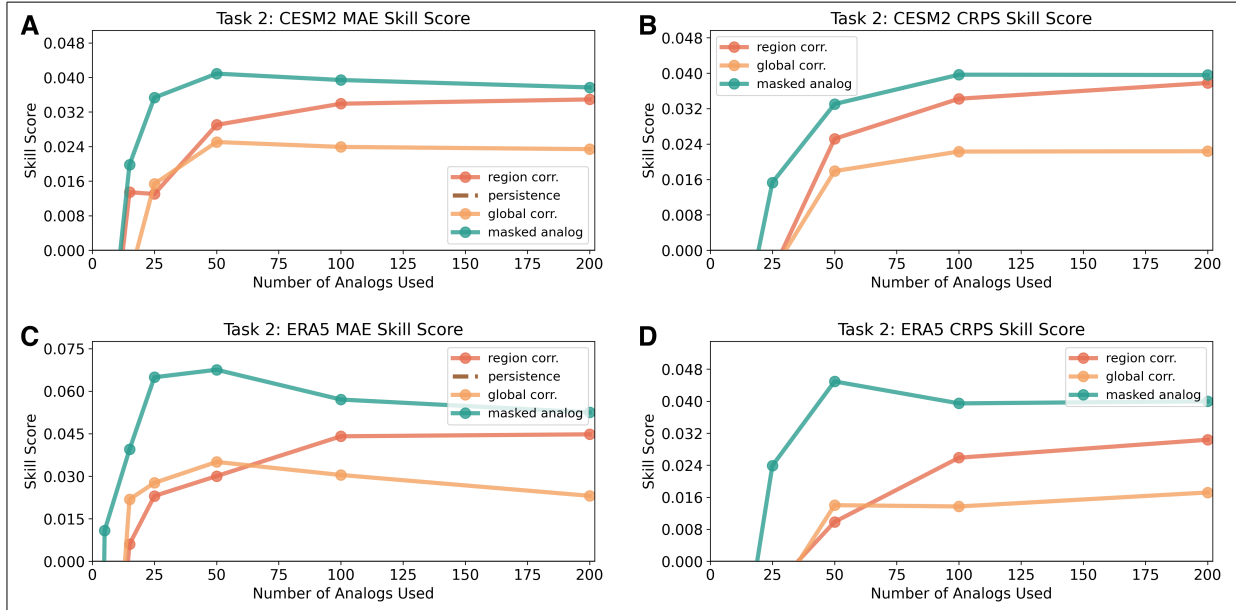


Figure 2.7: Skill scores for a) CESM2-LE accuracy, b) CESM2-LE CRPS, c) ERA5 accuracy, and d) ERA5 CRPS for Month 1 midwestern U.S. temperature regression.

extreme samples. This analysis highlights how the skill gap between AI-informed analogs and traditionally-selected analogs widens for more extreme temperature events. Moreover, as we use *predicted* extremity to discard samples, this information is available *a priori*, allowing forecasters to better understand when the analog ensemble is likely to perform best and building trust in its more extreme predictions. This behavior also holds for CESM2-LE data (Figure D.3).

2.3.3 Month 1-2 Mask Exploration in the North Atlantic

Lastly, we explore the learned mask’s ability to perform grid-point classification of upper atmospheric winds in the North Atlantic ($25^{\circ} - 48^{\circ}\text{N}$, $0^{\circ} - 80^{\circ}\text{W}$) and probe the mask itself to better understand the relative importance of different areas for successful prediction. At each grid point in the target region (rather than averaging across the target region), we classify the 250 hPa zonal wind (U250) using terciles. We make predictions for December-January and January-February, using the learned mask in Figure 2.9. In this case, we select analogs using both U250 and surface temperature as inputs. Therefore, we learn a unique mask for each field, although, importantly, these masks are learned together by the network.

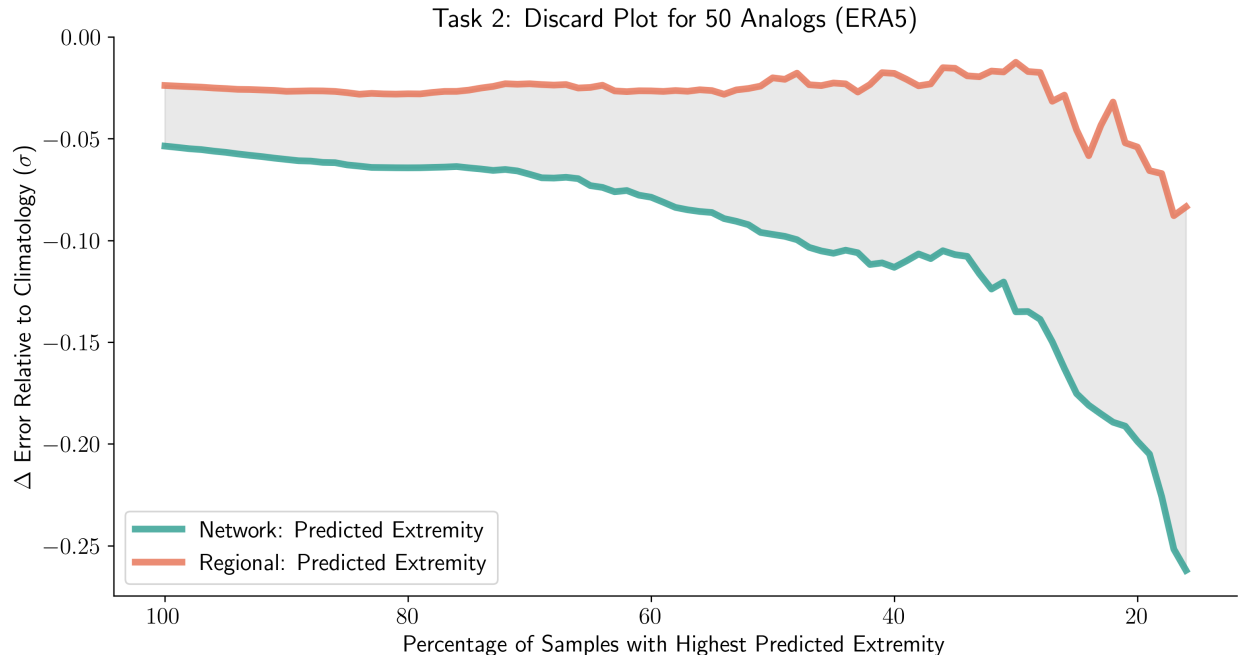


Figure 2.8: Discard plot based on predicted extremity with an ensemble of 50 analogs for midwestern U.S. summer temperature regression, testing on ERA5 data.

We evaluate skill at each grid point in the whole field (e.g., Figure D.4), summarizing these with mean skill scores over the target region (Figure 2.10). The learned mask outperforms all baselines for this grid-point-by-grid-point classification. Skill peaks at 400 analogs for both CESM2-LE and ERA5 data, except for CESM2-LE classification, which peaks at 800 analogs.

Additionally, we explore the learned masks to better understand the relative importance of initial conditions in different areas for successful analog prediction. We do so by ablating the mask, i.e. setting the weights to 0, and observing changes in skill. We analyze changes in BS skill with CESM2-LE data and a 400 analog ensemble. We test on CESM2-LE data rather than ERA5 because the impacts on BS are small, and thus, there are too few samples to draw meaningful conclusions from ERA5 data. We employ three ablation methods: 1) Ablating entire fields (e.g. temperature or U250), 2) threshold ablation, where we increase mask sparsity by setting weights to 0 if they are below either the 40th, 80th, or 90th percentile or incentivizing sparsity during training itself by adding constrained inverse L_2 regularization (see B.2 for details) and 3) ablating specific regions (e.g. the Northern Hemisphere). Masks for examples of these ablation methods are shown

in Figure 2.11. All of these ablation methods, except constrained inverse L_2 regularization, are performed after the mask has already been learned.

We focus on a 400 analog ensemble, as this is the number of analogs for which skill peaks (Figure 2.12), although the general trends remain similar across ensemble size (Figure D.6). We find a slight improvement in skill when we increase mask sparsity by thresholding or by introducing constrained inverse L_2 regularization. This increase in skill with a sparser map is consistent with [33], who found a slight improvement in skill for multi-year predictions using a $\sim 95\%$ percentile threshold. Though temperature appears to be the more important of the two fields, ablating either the temperature field or the U250 field results in a significant decrease in skill, highlighting the importance of both for identifying skillful analogs. While all ablation methods besides increasing sparsity decreases skill, ablating the Northern Hemisphere, both fields, and ocean temperatures result in the largest drop in skill.

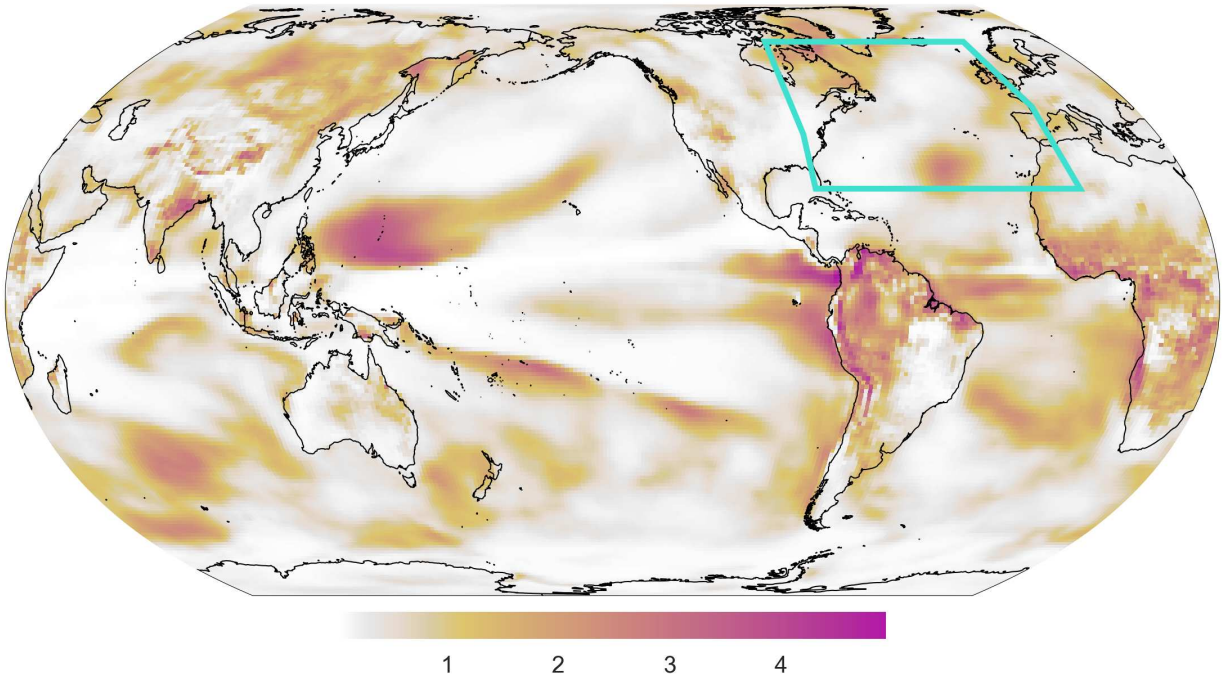
2.4 Conclusions

We demonstrate how an AI-informed model analog forecasting approach, previously only shown to be skillful on seasonal-to-decadal timescales, can also produce skillful deterministic and probabilistic subseasonal-to-seasonal predictions. We showcase this approach’s improvement over climatological, persistence, and traditional analog forecasting methods in both a perfect model framework and with reanalysis data for classification and regression tasks. Moreover, the interpretability of this AI-based approach enables users to explore the learned mask to gain insight into the relative importance of different areas of the globe for successful prediction. We perform an analysis of the learned mask for North Atlantic U250 classification, assessing which variable fields and areas of the globe are comparatively more important for a successful analog forecast. This type of analysis can help identify key initial conditions that most influence climate state evolution on S2S timescales, guiding both future model development and observational prioritization.

The AI-informed analog ensembles additionally provides both improved uncertainty quantification and prediction of extremes compared to traditional analog forecasting methods. Since ex-

treme temperature events have a disproportionate impact on human health, agriculture, and energy/water management, improving their prediction is essential for mitigating the most dire consequences [1]. Furthermore, accurate measures of forecast uncertainty are crucial for identifying windows of opportunity. By knowing in advance when a prediction is likely to be more or less skillful, forecasters can better gauge their confidence and decision-makers can more effectively assess risk [36]. Thus, with better uncertainty quantification, extreme prediction, and overall skill, the AI-informed analog approach offers a promising avenue for improving subseasonal-to-seasonal analog forecasts, with the potential to enhance decision-making across a variety of sectors.

Task 3: Temperature Mask for the North Atlantic



Task 3: U250 Mask for the North Atlantic

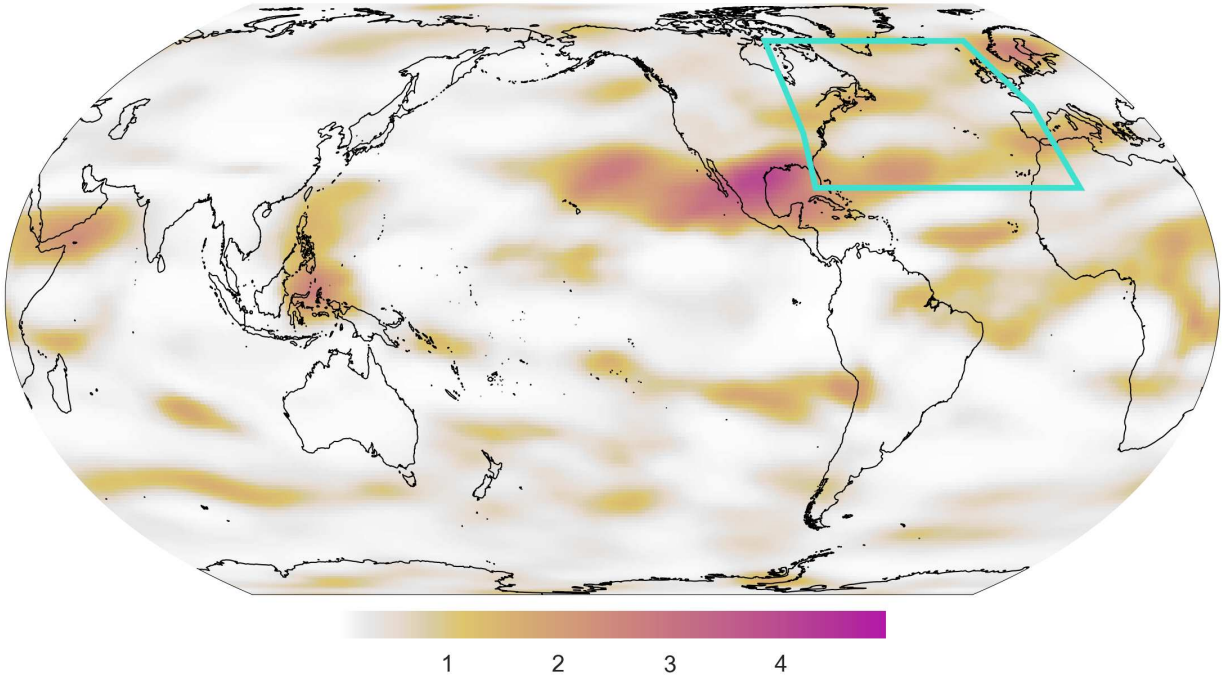


Figure 2.9: The learned mask for Task #3, North Atlantic winter U250 classification. The cyan box outlines the target region.

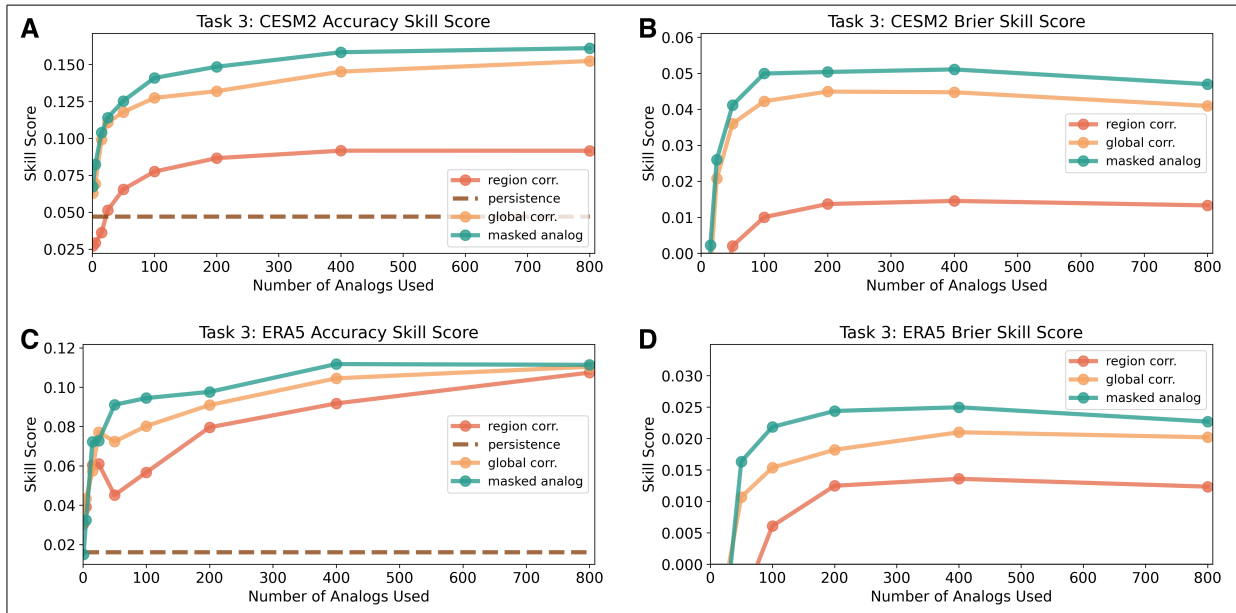


Figure 2.10: Skill scores for a) CSM2-LE accuracy, b) CSM2-LE BS, c) ERA5 accuracy, and d) ERA5 BS for Month 1-2 North Atlantic U250 classification.

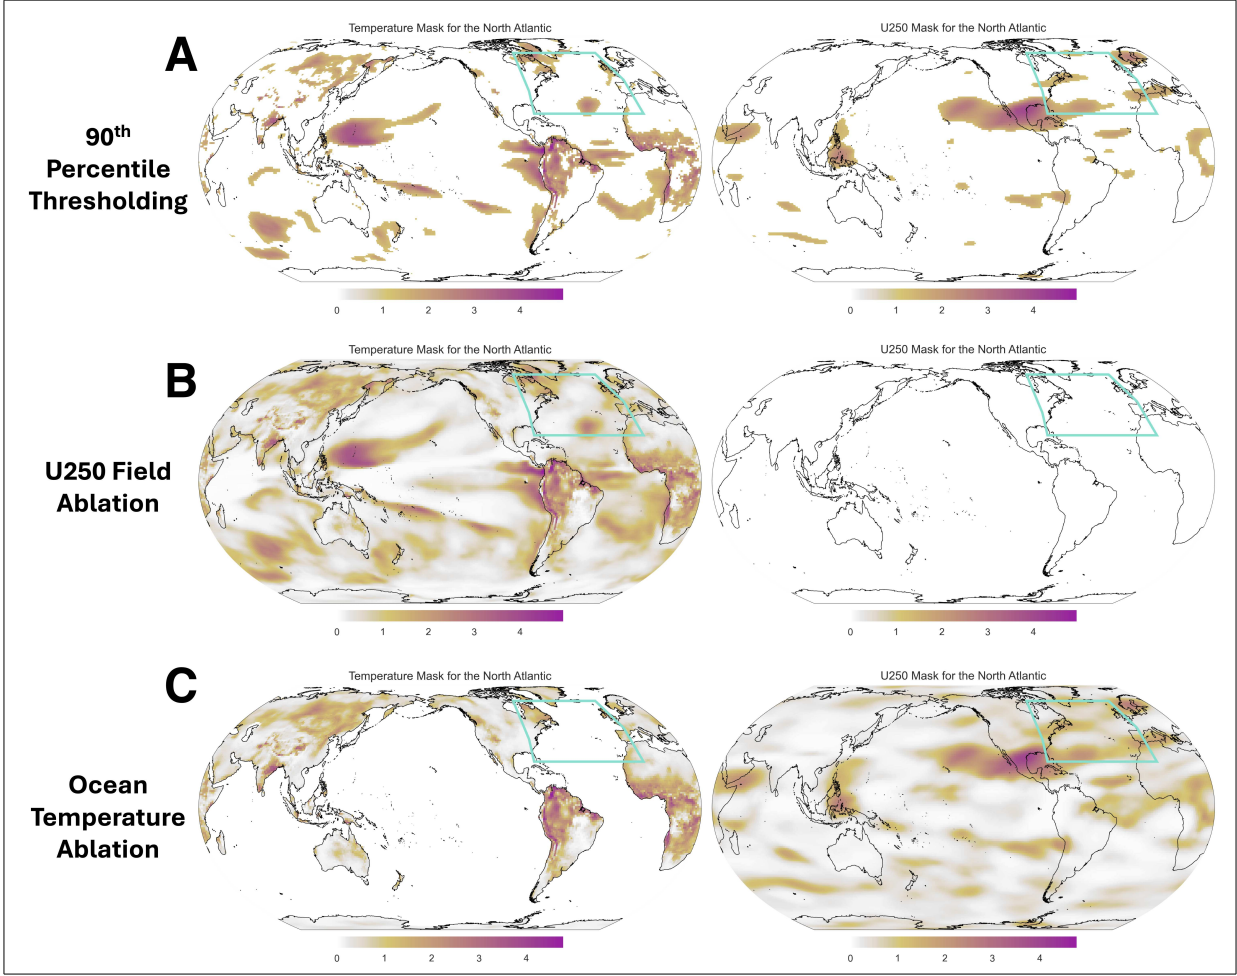


Figure 2.11: Examples of masks with each of the ablation methods: (a) Ablating entire fields, (b) Threshold ablation, and (c) Ablating specific regions.




Figure 2.12: BS skill for different ablation methods evaluated on CESM2-LE data.

Chapter 3

Discussion and Future Work

In this work, we achieve skillful results, while applying a similar framework to [33], where a single mask is learned for each prediction problem. However, another viable approach, and one used in [37], is to learn state-dependent masks (i.e., learning a unique mask for every SOI). For instance, one could pass each SOI through the neural network to generate a unique mask or one could train the mask on subsets of data (e.g., learning masks by training with subsets of SOIs that exhibit different phases of the MJO). We experimented with a network architecture that allowed the mask to be state-dependent, and passed each SOI through the network to generate unique masks. However, we found almost no variation in weights across the masks and no change in skill. While, this lack of meaningful variation may be because a single mask captures the most important features with little variation between SOIs, the network architecture itself could also have inhibited more disparate masks. Future work could explore this approach further, developing a network architecture that prioritizes learning state-dependent masks.

Like previous work [33, 37, 11], we also leverage a model analog approach rather than rely on observations alone to form our analog library. While we use a single model (CESM2) to form our analog library, previous work has experimented with using multiple models on multi-year-to-decadal timescales to both form the analog library and to train their neural network [11]. Including more climate models or, specifically for S2S timescales, extended-range AI or dynamical weather forecasts [18, 40] may similarly boost skill on these shorter timescales.

Further, although we have shown that this method produces skillful S2S predictions on near-observational data, the approach is still limited by the model's, in this case CESM2's, biases [30, 42, 45]. Future work could also explore how to best incorporate near-observational data into the AI-based analog forecasting approach. For example, [11] found that transfer learning from models to reanalysis improved analog prediction skill; a similar approach could be useful for learning masks on S2S timescales as well. Moreover, reanalysis data could also be incorporated into the

analog library itself, perhaps offering even more realistic climate trajectories than those from a library composed solely of model data. These additions to the analog library and/or training set could further advance AI-based analog S2S forecasts, especially for regions and seasons where model biases are more pronounced.

Bibliography

- [1] Amir AghaKouchak et al. “Climate Extremes and Compound Hazards in a Warming World”. In: *Annual Review of Earth and Planetary Sciences* 48. Volume 48, 2020 (2020), pp. 519–548. ISSN: 1545-4495. DOI: <https://doi.org/10.1146/annurev-earth-071719-055228>. URL: <https://www.annualreviews.org/content/journals/10.1146/annurev-earth-071719-055228>.
- [2] John R. Albers and Matthew Newman. “A Priori Identification of Skillful Extratropical Subseasonal Forecasts”. In: *Geophysical Research Letters* 46.21 (2025/01/09 2019), pp. 12527–12536. DOI: <https://doi.org/10.1029/2019GL085270>. URL: <https://doi.org/10.1029/2019GL085270>.
- [3] Omid Alizadeh. “Advances and challenges in climate modeling”. In: *Climatic Change* 170.1 (2022), p. 18. DOI: [10.1007/s10584-021-03298-4](https://doi.org/10.1007/s10584-021-03298-4). URL: <https://doi.org/10.1007/s10584-021-03298-4>.
- [4] Richard B. Alley, Kerry A. Emanuel, and Fuqing Zhang. “Advances in weather prediction”. In: *Science* 363.6425 (2025/04/18 2019), pp. 342–344. DOI: [10.1126/science.aav7274](https://doi.org/10.1126/science.aav7274). URL: <https://doi.org/10.1126/science.aav7274>.
- [5] Melissa L. Breeden et al. “The Spring Minimum in Subseasonal 2-m Temperature Forecast Skill over North America”. In: *Monthly Weather Review* 150.10 (2022), pp. 2617–2628. DOI: <https://doi.org/10.1175/MWR-D-22-0062.1>. URL: <https://journals.ametsoc.org/view/journals/mwre/150/10/MWR-D-22-0062.1.xml>.
- [6] Lei Chen et al. “A machine learning model that outperforms conventional global subseasonal forecast models”. In: *Nature Communications* 15.1 (2024), p. 6425. DOI: [10.1038/s41467-024-50714-1](https://doi.org/10.1038/s41467-024-50714-1). URL: <https://doi.org/10.1038/s41467-024-50714-1>.
- [7] G. Danabasoglu et al. “The Community Earth System Model Version 2 (CESM2)”. In: *Journal of Advances in Modeling Earth Systems* 12.2 (2025/01/17 2020), e2019MS001916. DOI: <https://doi.org/10.1029/2019MS001916>. URL: <https://doi.org/10.1029/2019MS001916>.

- [8] Hui Ding et al. “Diagnosing Secular Variations in Retrospective ENSO Seasonal Forecast Skill Using CMIP5 Model-Analogs”. In: *Geophysical Research Letters* 46.3 (2019/01/09 2019), pp. 1721–1730. DOI: <https://doi.org/10.1029/2018GL080598>. URL: <https://doi.org/10.1029/2018GL080598>.
- [9] Hui Ding et al. “Skillful Climate Forecasts of the Tropical Indo-Pacific Ocean Using Model-Analogs”. In: *Journal of Climate* 31.14 (2018), pp. 5437–5459. DOI: <https://doi.org/10.1175/JCLI-D-17-0661.1>. URL: <https://journals.ametsoc.org/view/journals/clim/31/14/jcli-d-17-0661.1.xml>.
- [10] Daniela I. V. Domeisen et al. “Advances in the Subseasonal Prediction of Extreme Events: Relevant Case Studies across the Globe”. In: *Bulletin of the American Meteorological Society* 103.6 (2022), E1473–E1501. DOI: <https://doi.org/10.1175/BAMS-D-20-0221.1>. URL: <https://journals.ametsoc.org/view/journals/bams/103/6/BAMS-D-20-0221.1.xml>.
- [11] M. A. Fernandez and Elizabeth A. Barnes. *Multi-Year-to-Decadal Temperature Prediction using a Machine Learning Model-Analog Framework*. 2025. arXiv: 2502.17583 [physics.aos-ph]. URL: <https://arxiv.org/abs/2502.17583>.
- [12] Laura Ferranti et al. “How far in advance can we predict changes in large-scale flow leading to severe cold conditions over Europe?” In: *Quarterly Journal of the Royal Meteorological Society* 144.715 (2018/03/17 2018), pp. 1788–1802. DOI: <https://doi.org/10.1002/qj.3341>. URL: <https://doi.org/10.1002/qj.3341>.
- [13] Mark C. Freeman, Ben Groom, and Richard J. Zeckhauser. “Better predictions, better allocations: scientific advances and adaptation to climate change”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 373.2055 (2015/03/31 2015), p. 20150122. DOI: 10.1098/rsta.2015.0122. URL: <https://doi.org/10.1098/rsta.2015.0122>.

- [14] Ji-Young Han et al. “Ensemble size versus bias correction effects in subseasonal-to-seasonal (S2S) forecasts”. In: *Geoscience Letters* 10.1 (2023), p. 37. DOI: 10.1186/s40562-023-00292-9. URL: <https://doi.org/10.1186/s40562-023-00292-9>.
- [15] Hans Hersbach et al. “The ERA5 global reanalysis”. In: *Quarterly Journal of the Royal Meteorological Society* 146.730 (2020/01/20 2020), pp. 1999–2049. DOI: <https://doi.org/10.1002/qj.3803>. URL: <https://doi.org/10.1002/qj.3803>.
- [16] T. N. Krishnamurti et al. “Improved Weather and Seasonal Climate Forecasts from Multimodel Superensemble”. In: *Science* 285.5433 (2020/02/11 1999), pp. 1548–1550. DOI: 10.1126/science.285.5433.1548. URL: <https://doi.org/10.1126/science.285.5433.1548>.
- [17] Daniel Kull et al. *The Value of Surface-based Meteorological Observation Data*. Jan. 2021. DOI: 10.1596/35178.
- [18] Simon Lang et al. *AIFS – ECMWF’s data-driven forecasting system*. 2024. arXiv: 2406.01465 [physics.aos-ph]. URL: <https://arxiv.org/abs/2406.01465>.
- [19] M. Leutbecher and T. N. Palmer. “Ensemble forecasting”. In: *Journal of Computational Physics* 227.7 (2008), pp. 3515–3539. DOI: <https://doi.org/10.1016/j.jcp.2007.02.014>. URL: <https://www.sciencedirect.com/science/article/pii/S0021999107000812>.
- [20] Fenghua Ling et al. *FengWu-W2S: A deep learning model for seamless weather-to-subseasonal forecast of global atmosphere*. 2024. arXiv: 2411.10191 [cs.LG]. URL: <https://arxiv.org/abs/2411.10191>.
- [21] Edward N. Lorenz. “Atmospheric Predictability as Revealed by Naturally Occurring Analogues”. In: *Journal of Atmospheric Sciences* 26.4 (1969), pp. 636–646. DOI: [https://doi.org/10.1175/1520-0469\(1969\)26<636:APARBN>2.0.CO;2](https://doi.org/10.1175/1520-0469(1969)26<636:APARBN>2.0.CO;2). URL: https://journals.ametsoc.org/view/journals/atsc/26/4/1520-0469_1969_26_636_aparbn_2_0_co_2.xml.
- [22] Jiale Lou, Matthew Newman, and Andrew Hoell. “Multi-decadal variation of ENSO forecast skill since the late 1800s”. In: *npj Climate and Atmospheric Science* 6.1 (2023), p. 89. DOI: 10.1038/s41612-023-00417-z. URL: <https://doi.org/10.1038/s41612-023-00417-z>.

- [23] R. Mahmood et al. “Constraining low-frequency variability in climate projections to predict climate on decadal to multi-decadal timescales – a poor man’s initialized prediction system”. In: *Earth System Dynamics* 13.4 (2022), pp. 1437–1450. DOI: 10.5194/esd-13-1437-2022. URL: <https://esd.copernicus.org/articles/13/1437/2022/>.
- [24] Annarita Mariotti et al. “Windows of Opportunity for Skillful Forecasts Subseasonal to Seasonal and Beyond”. In: *Bulletin of the American Meteorological Society* 101.5 (2020), E608–E625. DOI: <https://doi.org/10.1175/BAMS-D-18-0326.1>. URL: <https://journals.ametsoc.org/view/journals/bams/101/5/bams-d-18-0326.1.xml>.
- [25] Kirsten J. Mayer and Elizabeth A. Barnes. “Subseasonal Forecasts of Opportunity Identified by an Explainable Neural Network”. In: *Geophysical Research Letters* 48.10 (2025/01/09 2021), e2020GL092092. DOI: <https://doi.org/10.1029/2020GL092092>. URL: <https://doi.org/10.1029/2020GL092092>.
- [26] Patrick L. McDermott and Christopher K. Wikle. “A model-based approach for analog spatio-temporal dynamic forecasting”. In: *Environmetrics* 27.2 (2025/01/09 2016), pp. 70–82. DOI: <https://doi.org/10.1002/env.2374>. URL: <https://doi.org/10.1002/env.2374>.
- [27] William J. Merryfield et al. “Current and Emerging Developments in Subseasonal to Decadal Prediction”. In: *Bulletin of the American Meteorological Society* 101.6 (2020), E869–E896. DOI: <https://doi.org/10.1175/BAMS-D-19-0037.1>. URL: <https://journals.ametsoc.org/view/journals/bams/101/6/bamsD190037.xml>.
- [28] A. Brett Mullan and Craig S. Thompson. “Analogue forecasting of New Zealand climate anomalies”. In: *International Journal of Climatology* 26.4 (2025/01/09 2006), pp. 485–504. DOI: <https://doi.org/10.1002/joc.1261>. URL: <https://doi.org/10.1002/joc.1261>.
- [29] T. N. Palmer et al. “DEVELOPMENT OF A EUROPEAN MULTIMODEL ENSEMBLE SYSTEM FOR SEASONAL-TO-INTERANNUAL PREDICTION (DEMETER)”. In: *Bulletin of the American Meteorological Society* 85.6 (2004), pp. 853–872. DOI: <https://doi.org/10.1175/JCLI1500853>.

- org/10.1175/BAMS-85-6-853. URL: <https://journals.ametsoc.org/view/journals/bams/85/6/bams-85-6-853.xml>.
- [30] Da Pang, Xianghui Fang, and Lei Wang. “Feedback processes responsible for the deficiency of El Niño diversity in CESM2”. In: *Climate Dynamics* 63.1 (2024), p. 47. DOI: 10.1007/s00382-024-07515-5. URL: <https://doi.org/10.1007/s00382-024-07515-5>.
- [31] Kathy Pegion et al. “The Subseasonal Experiment (SubX): A Multimodel Subseasonal Prediction Experiment”. In: *Bulletin of the American Meteorological Society* 100.10 (2019), pp. 2043–2060. DOI: <https://doi.org/10.1175/BAMS-D-18-0270.1>. URL: <https://journals.ametsoc.org/view/journals/bams/100/10/bams-d-18-0270.1.xml>.
- [32] Yihao Peng et al. “Skill improvement of the yearly updated reforecasts in ECMWF S2S prediction from 2016 to 2022”. In: *Atmospheric and Oceanic Science Letters* 16.5 (2023), p. 100357. DOI: <https://doi.org/10.1016/j.aosl.2023.100357>. URL: <https://www.sciencedirect.com/science/article/pii/S1674283423000351>.
- [33] Jamin K. Rader and Elizabeth A. Barnes. “Optimizing Seasonal-To-Decadal Analog Forecasts With a Learned Spatially-Weighted Mask”. In: *Geophysical Research Letters* 50.23 (2025/01/09 2023), e2023GL104983. DOI: <https://doi.org/10.1029/2023GL104983>. URL: <https://doi.org/10.1029/2023GL104983>.
- [34] Andrew W. Robertson et al. “Summary of workshop on sub-seasonal to seasonal predictability of extreme weather and climate”. In: *npj Climate and Atmospheric Science* 1.1 (2018), p. 20178. DOI: 10.1038/s41612-017-0009-1. URL: <https://doi.org/10.1038/s41612-017-0009-1>.
- [35] Isla R. Simpson et al. “The CESM2 Single-Forcing Large Ensemble and Comparison to CESM1: Implications for Experimental Design”. In: *Journal of Climate* 36.17 (2023), pp. 5687–5711. DOI: <https://doi.org/10.1175/JCLI-D-22-0666.1>. URL: <https://journals.ametsoc.org/view/journals/clim/36/17/JCLI-D-22-0666.1.xml>.

- [36] Julia Slingo and Tim Palmer. “Uncertainty in Weather and Climate Prediction”. In: *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences* 369 (Dec. 2011), pp. 4751–67. DOI: 10.1098/rsta.2011.0161.
- [37] Kinya Toride et al. *Using Deep Learning to Identify Initial Error Sensitivity for Interpretable ENSO Forecasts*. 2024. arXiv: 2404.15419 [physics.ao-ph]. URL: <https://arxiv.org/abs/2404.15419>.
- [38] H. M. Van den Dool. “Searching for analogues, how long must we wait?” In: *Tellus A* 46.3 (2025/01/09 1994), pp. 314–324. DOI: <https://doi.org/10.1034/j.1600-0870.1994.t01-2-00006.x>. URL: <https://doi.org/10.1034/j.1600-0870.1994.t01-2-00006.x>.
- [39] Frédéric Vitart and Andrew W. Robertson. “The sub-seasonal to seasonal prediction project (S2S) and the prediction of extreme events”. In: *npj Climate and Atmospheric Science* 1.1 (2018), p. 3. DOI: 10.1038/s41612-018-0013-0. URL: <https://doi.org/10.1038/s41612-018-0013-0>.
- [40] Frédéric Vitart et al. “The next extended-range configuration for IFS Cycle 48r1”. In: (173 Oct. 2022), pp. 21–26. DOI: 10.21957/fv6k37c49h. URL: <https://www.ecmwf.int/node/20521>.
- [41] John Walsh et al. “An Analog Method for Seasonal Forecasting in Northern High Latitudes”. In: *Atmospheric and Climate Sciences* 11 (Jan. 2021), pp. 469–485. DOI: 10.4236/acs.2021.113028.
- [42] Ho-Hsuan Wei et al. “Tropical Pacific Air-Sea Interaction Processes and Biases in CESM2 and Their Relation to El Niño Development”. In: *Journal of Geophysical Research: Oceans* 126.6 (2025/03/04 2021), e2020JC016967. DOI: <https://doi.org/10.1029/2020JC016967>. URL: <https://doi.org/10.1029/2020JC016967>.
- [43] A. Weisheimer and Tim Palmer. “On the reliability of seasonal climate forecasts”. In: *Journal of the Royal Society, Interface / the Royal Society* 11 (July 2014), p. 20131162. DOI: 10.1098/rsif.2013.1162.

- [44] Christopher J. White et al. “Potential applications of subseasonal-to-seasonal (S2S) predictions”. In: *Meteorological Applications* 24.3 (2017), pp. 315–325. DOI: <https://doi.org/10.1002/met.1654>. eprint: <https://rmets.onlinelibrary.wiley.com/doi/pdf/10.1002/met.1654>. URL: <https://rmets.onlinelibrary.wiley.com/doi/abs/10.1002/met.1654>.
- [45] M. D. Woelfle et al. “Evolution of the Double-ITCZ Bias Through CESM2 Development”. In: *Journal of Advances in Modeling Earth Systems* 11.7 (2025/03/04 2019), pp. 1873–1893. DOI: <https://doi.org/10.1029/2019MS001647>. URL: <https://doi.org/10.1029/2019MS001647>.
- [46] Yanling Wu and Xiaoqin Yan. “Evaluating Changes in the Multiyear Predictability of the Pacific Decadal Oscillation Using Model Analogs since 1900”. In: *Journal of Marine Science and Engineering* 11.5 (2023). DOI: 10.3390/jmse11050980.
- [47] Chidong Zhang. “Madden–Julian Oscillation: Bridging Weather and Climate”. In: *Bulletin of the American Meteorological Society* 94.12 (2013), pp. 1849–1870. DOI: <https://doi.org/10.1175/BAMS-D-12-00026.1>. URL: <https://journals.ametsoc.org/view/journals/bams/94/12/bams-d-12-00026.1.xml>.
- [48] Zhizhen Zhao and Dimitrios Giannakis. “Analog Forecasting with Dynamics-Adapted Kernels”. In: *Nonlinearity* 29 (Aug. 2016). DOI: 10.1088/0951-7715/29/9/2888.

Appendix A

Data and Code Availability

CESM2-LE data is available at <https://www.earthsystemgrid.org/dataset/ucar.cgd.cesm2le.output.html> [7]. ERA5 data is available at <https://cds.climate.copernicus.eu/datasets> [15]. Code is available at <https://github.com/jlandsbe/S2S>.

Appendix B

Neural Network Details

B.1 Hyperparameters

We use the following parameters for the neural network:

Table B.1: Hyperparameters used in the model training.

Parameter	Value
Optimizer	Adam
Learning rate	.0001
Batch size	32
Loss function	Mean squared error
Validation batch size	1000
Early stopping patience	30
Early stopping minimum delta	0.0001

B.2 Constrained Inverse L_2 Regularization

To increase mask sparsity, we implement constrained inverse L_2 regularization. We do so to compare how post-hoc thresholding and learned sparsity compare in terms of model performance.

We add the following term to the loss function:

$$\frac{\lambda_2}{\sqrt{\sum_{i=1}^n w_i^2}} \tag{B.1}$$

where λ_2 is the regularization strength, w_i is the weight of the i th grid point, and n is the total number of grid points in the weighted mask.

This term is restricted during training such that $\sum_{i=1}^n w_i = n$. With this constraint, the regularization term is maximized (high loss) when $\forall i, w_i = 1$ and minimized when, for some $j, w_j = n$ and $w_{i \neq j} = 0$. Thus, this term promotes having more disparate weight values, with some weights

of very high values and some weights of very low values. We generate the mask in Figure D.5 for the North Atlantic (Task #3), using this regularization term with $\lambda_2 = 100$.

Appendix C

Data: CESM2-LE Members

For monthly CESM2-LE data, we use the following ensemble members:

Table C.1: List of monthly ensemble members used for the analog library, training SOIs, validation SOIs, and testing SOIs.

Monthly Member Type	Members
Analog Library Members	1301.020, 1301.019, 1301.018, 1301.017, 1301.016, 1301.015, 1301.014, 1301.013, 1301.012, 1301.011, 1281.020, 1281.019, 1281.018, 1281.017, 1281.016, 1281.015, 1281.014, 1281.013, 1281.012
SOI Train Members	1301.010, 1301.009, 1301.008, 1301.007, 1301.006, 1301.005, 1301.004, 1301.003, 1301.002, 1301.001
SOI Validation Members	1281.010, 1281.009
SOI Test Members	1281.001, 1281.002

For daily CESM2-LE data, we use the following ensemble members:

Table C.2: List of daily ensemble members used for the analog library, training SOIs, validation SOIs, and testing SOIs.

Daily Member Type	Members
Analog Library Members	1231.012, 1251.013, 1251.014, 1281.015, 1281.016
SOI Train Members	1301.017, 1301.018
SOI Validation Members	1281.017
SOI Test Members	1251.019

Appendix D

Figures




Figure D.1: Discard plot for Week 3-4 Southern California temperature classification using 1500 analogs, testing on CESM2-LE data. Unlike when testing on reanalysis data, global ensemble agreement serves as a reasonable metric for forecast uncertainty. However, the learned-mask ensemble agreement still showcases a greater decrease in error for states with higher ensemble agreement.




Figure D.2: With regression problems, as the number of analogs approaches the total library size, the prediction becomes more and more similar to a climatological prediction. In this example the library size is ~ 10000 , where the skill is 0.

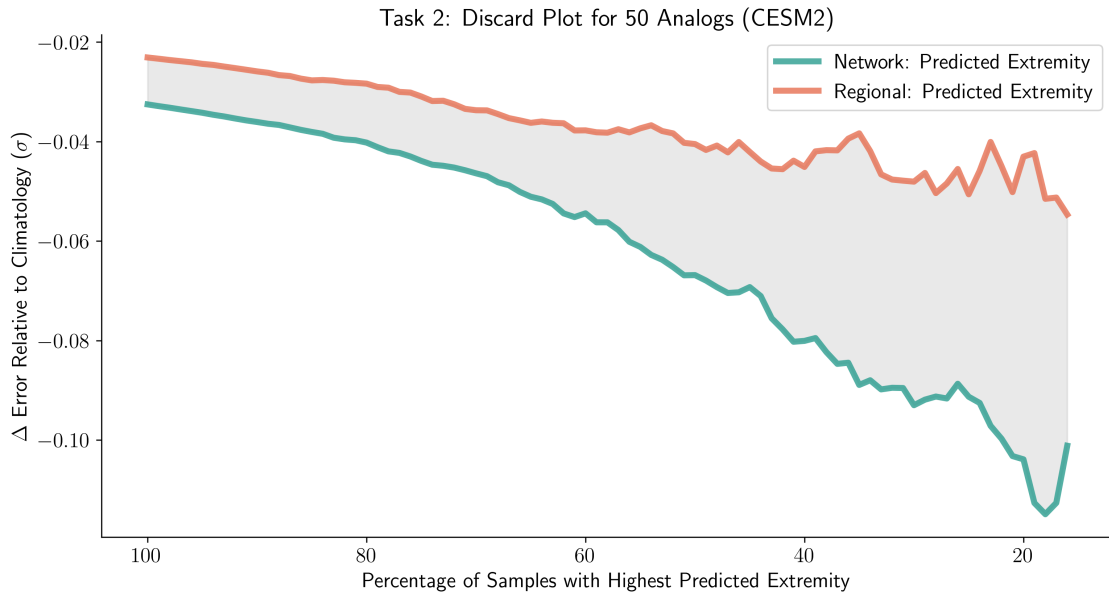


Figure D.3: A discard plot for an ensemble of 50 analogs for midwestern U.S. summer temperature regression, testing on CESM2-LE data. While the regional mask perform relatively better on extreme predictions for CESM2-LE compared to ERA5 data, there is still a much larger improvement for extreme predictions with the learned mask.

Task 3: ERA5 MAE Skill Relative to Climatology with 400 analogs

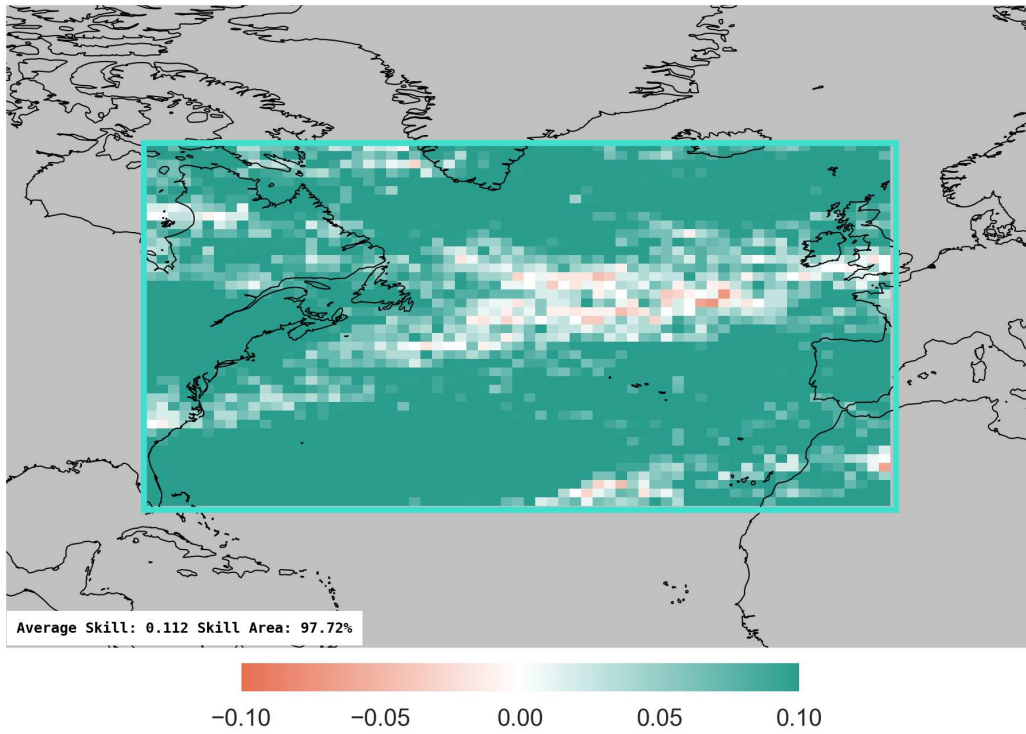
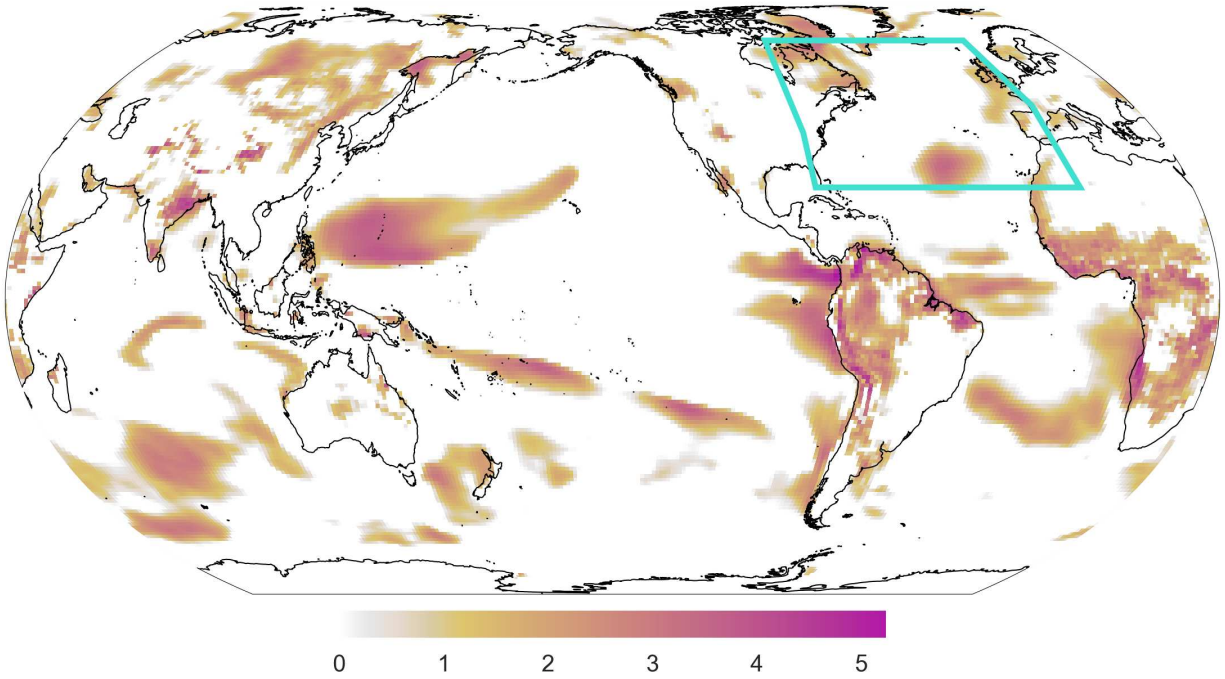


Figure D.4: In Task #3 a field of values rather than the average value across the region of interest is predicted. This map shows the average skill across SOIs when using a 400 analog ensemble on ERA5 data. At almost every grid point, the learned mask outperforms the climatological prediction.

Task 3: Temperature Mask for the North Atlantic



Task 3: U250 Mask for the North Atlantic

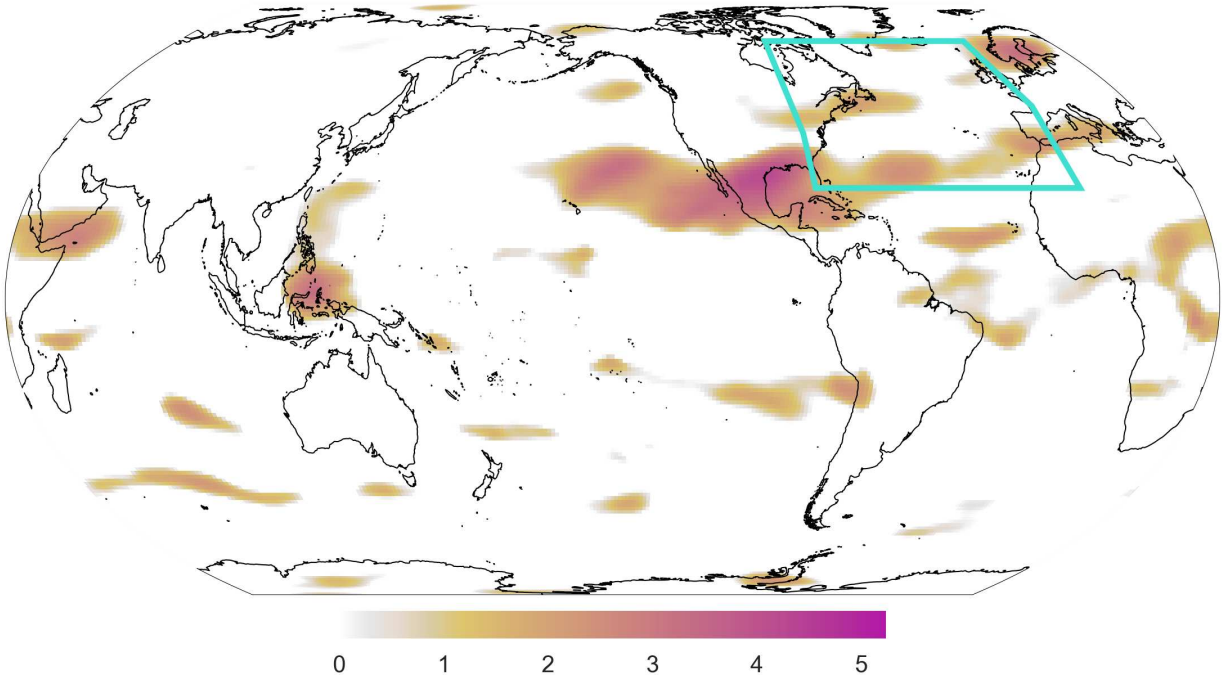


Figure D.5: Using inverse L_2 regularization with $\lambda_2 = 100$ for the North Atlantic (Task #3) results in a sparser map.




Figure D.6: Expanded table, showing changes in BS skill with different ablation methods, for 100-800 analogs