Noname manuscript No.
(will be inserted by the editor)

Predicting regional and pan-Arctic sea ice anomalies with kernel analog forecasting

Darin Comeau · Dimitrios Giannakis · Zhizhen Zhao · Andrew J. Majda

Received: date / Accepted: date

Abstract Predicting the Arctic sea ice extent is a notoriously difficult forecasting problem, even from lead times as short as one month. Motivated by Arctic intra-annual variability phenomena such as sea surface temperature reemergence and sea ice reemergence, we use a prediction approach for sea ice anomalies based on analog forecasting. Traditional analog forecasting relies on identifying a single analog in a historical record, usually by minimizing Euclidean distance, and forming a forecast from the analog's historical trajectory. We use an ensemble of analogs for our forecasts, where the ensemble weights are determined by a dynamics-adapted similarity kernel, which takes into account the nonlinear geometry on the underlying data manifold. We apply this method for forecasting regional and pan-Arctic sea ice concentration and volume anomalies from multi-century climate model data, and in many cases find improvement over the persistence forecast. Moreover the patterns of predictive skill we see by region and season are consistent with different types of sea ice reemergence.

D. Comeau
Center for Atmosphere Ocean Science, Courant Institute of Mathematical Sciences, New York University, New York, NY.
E-mail: dcomeau@lanl.gov *Present address:* of D. Comeau:
Climate, Ocean, Sea Ice Modeling Group, Computational Physics and Methods (CCS-2), Los Alamos National Laboratory, Los Alamos, NM.

D. Giannakis
Center for Atmosphere Ocean Science, Courant Institute of Mathematical Sciences, New York University, New York, NY.

Z. Zhao
Center for Atmosphere Ocean Science, Courant Institute of Mathematical Sciences, New York University, New York, NY.
Present address: of Z. Zhao:
Department of Electrical and Computer Engineering, Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL.

A. Majda
Center for Atmosphere Ocean Science, Courant Institute of Mathematical Sciences, New York University, New York, NY.

1 Introduction

Predicting the climate state of the Arctic, particularly with regards to sea ice extent, has been a subject of increased recent interest in part driven by record-breaking minimums in September sea ice extent in 2007 and again in 2012. As new areas of the Arctic become accessible, this has increasingly become an important practical problem in addition to a scientific one, e.g. navigating shipping routes (Smith and Stephenson, 2013). Many different approaches have been used recently to address Arctic sea ice prediction, including statistical frameworks (Lindsay et al, 2008; Wang et al, 2015), through inherent predictability within general circulation models (GCMs) (Blanchard-Wrigglesworth et al, 2011a,b; Chevallier et al, 2013; Tietsche et al, 2013, 2014; Day et al, 2014), and using dynamical models to predict observations (Zhang et al, 2008; Sigmond et al, 2013; Wang et al, 2013). These methods have varying degrees of success in predicting pan-Arctic and regional sea ice area or extent (area with at least 15% sea ice concentration) and to a lesser degree sea ice volume. Indeed, in sea ice prediction, current generation numerical models and data assimilation systems have little additional skill beyond simple persistence or damped persistence forecasts.

Following the 2007 September sea ice extent minimum, Study of Environmental Arctic Change (SEARCH) began soliciting forecasts of September sea ice extent for the Sea Ice Outlook (SIO) in an effort to improve operational forecasts, which since 2013 has been handled by the Sea Ice Prediction Network (SIPN). They have found that year to year variability, rather than methods, dominate the ensemble's success, and that extreme years are in general less predictable (Stroeve et al, 2014). The forecasts, given at one to three lead month times, had particular difficulty with the September 2012 (extreme low) and September 2013 (extreme high) sea ice extents. A more recent study of SIO model forecasts by Blanchard-Wrigglesworth et al (2015) highlighted the importance of initial conditions on predictability by performing an initial condition perturbation experiment and finding wide spread among models' response.

Accurately predicting aspects of Arctic sea ice is made difficult by a number of factors (Stroeve et al, 2014; Guemas et al, 2014). In particular, the challenge of the changing mean Arctic state can be viewed as separating the role of forced (external) and natural (internal) variability on predictability of the sea ice system. As the mean state of the Arctic is changing, variability of observed sea ice area has increased, in part due to thinner ice being more variable (in both thickness and extent), which is then harder to predict (Blanchard-Wrigglesworth et al, 2011a; Holland et al, 2011, 2008). Since the satellite record, all months have a downward trend in sea ice extent, the largest being for September (Stroeve et al, 2012). Moreover, as thicker multiyear ice is replaced by thinner, younger ice, the trends steepen (Stroeve et al, 2012). Ice thickness data is seen offering key predictive information for sea ice area / extent (Bushuk et al, 2017; Blanchard-Wrigglesworth and Bitz, 2014; Chevallier and Salas-Mélia, 2012; Lindsay et al, 2008; Wang et al, 2013), but such observational data sets do not yet exist in uniform spatial and temporal coverage (although it should be noted that the Pan-Arctic Ice Ocean Modeling and Assimilation System (Zhang and Rothrock, 2003) produces sea ice volume data by assimilating observations of sea ice concentration with a sea ice thickness model). Ice age, in particular area of ice of a certain age, is also seen as an important predictor, also of which there is no reliable observational record (Stroeve et al, 2012). For these reasons, the changing Arctic mean state complicates statistical predictions based on historical relationships (Holland and Stroeve, 2011; Stroeve et al, 2014).

There are both dynamic and thermodynamic elements that factor into sea ice predictability. Ice thickness predictability in the Arctic is dominated by dynamic, rather than thermodynamic properties (Blanchard-Wrigglesworth and Bitz, 2014; Tietsche et al, 2014). On the other hand, limits on September sea ice extent are primarily thermodynamic (related to amount of open water formation in melt season), whereas dynamic induced anomalies have smaller influence, except in a thin ice regime (Holland et al, 2011). Improvement in melt-pond parameterizations in the sea ice model Community Ice CodE (CICE) (Holland et al, 2012) have yielded skill in predicting September sea ice extent (Schröder et al, 2014), demonstrating potential predictive yield in improving process models.

Chaotic atmosphere variability also places an inherent limit of sea ice predictability (Day et al, 2014; Holland et al, 2011; Blanchard-Wrigglesworth et al, 2011b; Ogi et al, 2010) through its redistribution of sea ice. It has also been suggested that this importance may be lessened in a thinning sea ice regime, as historically high correlation between Arctic Oscillation and summer ice extent have been seen to weaken in recent years (Holland and Stroeve, 2011). Yet other studies (Stroeve et al, 2014) suggest that the importance of summer atmospheric conditions outweigh sea ice thickness in terms of providing predictive skill. The ocean temperature at depth has also been found to be an important predictor factor (Lindsay et al, 2008).

The problem of sea ice prediction becomes both of more practical use, while becoming more difficult, as we move from the pan-Arctic to regional scale, where local ice advection across regional boundaries and small scale influences on sea ice processes become important (Blanchard-Wrigglesworth and Bitz, 2014). Certain regions have been found to be more predictable than others; e.g. basins adjacent to Atlantic (Labrador to Barents) are more predictable than central Arctic basins (Day et al, 2014; Lindsay et al, 2008; Koenigk and Mikolajewicz, 2009), and the central Arctic basins that typically exhibit perennial sea ice cover are more difficult to predict than regions in the marginal ice zone (Day et al, 2014). In addition to the September sea ice extent metric, there has been increased focus on predicting regional sea ice advance and retreat dates (e.g. Sigmond et al (????)), and are now included as part of the SIO solicitation.

Sea ice reemergence is a phenomena where anomalies at one time reappear several months later, made evident by high lagged correlations, and has been found in both models and observations (Blanchard-Wrigglesworth et al, 2011a). Reemergence phenomena fall into two categories; one where anomalies from a melt season reemerge in the subsequent growth season, typically found in marginal ice zones, and are governed by ocean and large-scale atmospheric conditions, and another where anomalies from a growth season reemerge in the subsequent melt season, typically found in central Arctic regions that exhibit perennial sea ice, and are driven by sea ice thickness (Blanchard-Wrigglesworth et al, 2011a; Bushuk et al, 2014; Bushuk and Giannakis, 2015; Bushuk et al, 2015; Bushuk and Giannakis, 2017). This observed phenomena provides a promising source of sea ice predictability.

The timescales of predictability vary across studies, depending on the measure of predictive skill and the target month of prediction (among other factors), but generally fall in the 3–6 month range. While Lindsay et al (2008) found that most predictive information in the ice-ocean system is lost for lead times greater than 3 months, Blanchard-Wrigglesworth et al (2011a) found pan-Arctic sea ice area predictable for 1–2 years, and sea ice volume up to 3–4 years, in a perfect model framework. It has been found that predicting the state of sea ice in the spring is particularly difficult, with most of the predictive skill coming from fall persistence (Wang et al, 2013; Holland et al, 2011), and that March sea ice extent is largely uncorrelated with the following September sea ice extent (detrended) (Blanchard-Wrigglesworth et al, 2011a; Stroeve et al, 2014). While Day et al (2014) found a melt season 'predictability barrier', they also found that sea ice reemergence phenomena can aid in predictive skill, and this result was robust in their analysis of five GCMs.

Analog forecasting is an idea dating back to Lorenz (1969), where a prediction is made by identifying an appropriate historical analog to a given initial state, and using the analog's trajectory in the historical record to make a forecast of the present state. While this is attractive as a fully non-parametric, data-driven approach, a drawback of traditional analog forecasting is that it relies upon a single analog, usually identified by Euclidean distance, possibly introducing highly discontinuous behavior into the forecasting scheme. This can be improved upon by selecting an ensemble of analogs, and taking a weighted average of the associated trajectories. Analog forecasting has been used in numerous climate applications (Drosdowsky, 1994; Xavier and Goswami, 2007; Alessandrini et al, 2015), the latter of which also employed an ensemble approach. Given there are sources of sea ice predictability from the ocean, atmosphere, and sea ice itself (Guemas et al, 2014), a data-driven approach such as analog forecasting may be able to exploit complex coupled-system dynamics encoded in GCM data and provide skill in such a prediction problem.

In Zhao and Giannakis (2016) this idea was extended upon by assigning ensemble weights derived from a dynamics-adapted kernel, constructed in such a way as to give preferential weight to states with similar dynamics, referred to as kernel analog forecasting (KAF). Modes of variability intrinsic to the data analysis, as eigenfunctions of the kernel operator, are extracted with clean timescale separation and inherent predictability, while also being physically meaningful. KAF has been used in forecasting modes representing the Pacific Decadal Oscillation (PDO) and North Pacific Gyre Oscillation (NPGO) (Zhao and Giannakis, 2016), in which cases it was shown to be more skillful than parametric regression forecasting methods (Comeau et al, 2017). More recently KAF has been used in forecasting variability in the tropics by the Madden-Julian oscillation and the boreal summer intraseasonal oscillation (Alexander et al, 2017).

While KAF exhibits predictive skill in these extracted modes intrinsic to the data analysis, it is also desirable to have skill in forecasting objective observables that are independent of the analysis approach, e.g. Arctic sea ice anomalies (Comeau et al, 2017). The aim of this study is to extend upon Comeau et al (2017) by using KAF to study predictability of Arctic sea ice anomalies on various spatial and temporal scales in order to identify where and when we may (or may not) have predictability in this metric. Since utility of KAF depends upon the availability of an appropriately rich historical record, we examine predictability in a perfect model scenario, so there will be only natural variability present, with no external forced variability. Specifically, the aims of this study are:

1. *Spatial Impact on Predictability*: We consider various Arctic regions in the marginal ice zone, perennial ice zones, as well as the Arctic as a whole. The specific regions considered are detailed in Sect. 3.
2. *Temporal Impact on Predictability*: We break down error metrics by the target month of prediction to study seasonal effects. In particular, how well can we do in predicting the Arctic September sea ice extent anomaly?
3. *Predictor variables*: The predictor variables we consider are sea ice concentration (SIC), sea surface temperature (SST), sea level pressure (SLP), and sea ice thickness (SIT) data in order to gauge impact on predictive skill. Most of our analysis will not include SIT data due to its general unavailability as observational data.
4. *Predicting unobserved quantities*: As a strong test for the prediction methods, we aim to predict an unobserved quantity by targeting sea ice volume as our quantity to predict, without using SIT as a predictor variable.

The rest of this paper is structured as follows: The KAF method is described in Sect. 2. The data and experimental setup is described in Sect. 3, with the associated results in Sect. 4. Discussion and concluding remarks are given in Sect. 5.

2 Methods

The KAF method (Zhao and Giannakis, 2016; Comeau et al, 2017; Alexander et al, 2017), is designed to address the difficult task of prediction using massive data sets sampled from a complex nonlinear dynamical system in a very large state space. The motivating idea is to encode information from the underlying dynamics of the system into a kernel function, an exponentially decaying pairwise measure of similarity that can be loosely thought of a local covariance operator on the underlying data manifold. At the outset, during the training phase we have access to a time-ordered training data set $\{x_1,\ldots,x_n\}$ and the corresponding values $\{f_1,\ldots,f_n\}$ of a prediction observable. In our applications, the target observable is the aggregate sea ice anomaly over some region, and the training data are gridded climate variables. The main steps in KAF, outlined in detail below, are 1) perform Takens embedding of the data, 2) evaluate a dynamics-adapted similarity kernel on the embedded data, and 3) use weights from this kernel to make a forecast of an observable via out-of-sample extension formed by a weighted iterated sum.

2.1 Takens embedding

The first step in our analysis is to construct a new state variable through time-lagged embedding. For an embedding window of length q, which will depend on the time scale of our observable of interest, and a spatiotemporal series $z_1, z_2, \ldots z_n$ with $z_i \in \mathbb{R}^d$ (time index i), we form data set of x_i in lagged-embedded space (also called Takens embedding space) by

$$x_i = \left(z_i, z_{i-1}, \ldots, z_{i-(q-1)}\right) \in \mathbb{R}^{dq}.$$

The utility of this embedding is that it recovers the topology of the attractor of the underlying dynamical system through partial observations (the z_is) (Packard et al, 1980; Takens, 1981; Broomhead and King, 1986; Sauer et al, 1991; Robinson, 2005; Deyle and Sugihara, 2011). The choice of the embedding window q should be chosen long enough to capture the time-scales of interest, but not so long as to reduce the discriminating power of the kernel in determining locality.

2.2 Dynamics-adapted kernels

The kernel function we use to endow a geometry on our data manifold is from the Nonlinear Laplacian Spectral Analysis (NLSA) algorithm (Giannakis and Majda, 2012a,b, 2013, 2014). The kernel incorporates additional dynamic information by using phase velocities $\xi_i = \|x_i - x_{i-1}\|$, thus giving higher weight to regions of data space where the data is changing rapidly (see Giannakis (2015) for a geometrical interpretation), and is:

$$k\left(x_i, x_j\right) = \exp\left(-\frac{\|x_i - x_j\|^2}{\varepsilon \xi_i \xi_j}\right),$$

where ε is a scale parameter. We modify this to include multiple variables $x_i = \left(x_i^{(1)}, x_i^{(2)}\right)$ (Bushuk et al, 2014), possibly of different physical units, embedding windows, or grid points, which for two variables is

$$k\left(x_i, x_j\right) = \exp\left(-\frac{\|x_i^{(1)} - x_j^{(1)}\|^2}{\varepsilon \xi_i^{(1)} \xi_j^{(1)}} - \frac{\|x_i^{(2)} - x_j^{(2)}\|^2}{\varepsilon \xi_i^{(2)} \xi_j^{(2)}}\right), \qquad (1)$$

and extended to more than two variables in a similar manner. We next form row-normalized kernels,

$$P(x_i, x_j) = \frac{k(x_i, x_j)}{\sum_l k(x_i, x_l)}, \qquad (2)$$

which forms a row-stochastic matrix P that allows us to interpret each row as an empirical probability distribution of the second argument.

2.3 Out-of-sample extension via Laplacian pyramids

Our approach of assigning a value for a function f defined on a training data set X to a new test value $y \notin X$ will be through an out-of-sample extension technique known as Laplacian pyramids (Rabin and Coifman, 2012). In our context, the training data will be a spatio-temporal data set comprised of (lagged-embedded) state vectors x_i of gridded state variables (usually SIC, SST, and SLP), $f_i = f(x_i)$ is the function that gives us the sea ice cover anomaly of the state x_i, and y will be a new state vector (in lagged-embedded space), from which we would like to make a forecast of future sea ice cover anomalies.

We define a family of kernels P_l by modifying the NLSA kernels k in Eq. (1) to have scale parameter $\sigma_0/2^l$ rather than ε, which we denote k_l, and then P_l is the row-sum normalized k_l, as in Eq. (2). This forms a multiscale family of kernels with increasing dyadic resolution in l. A function f is approximated in a multiscale manner as an iterated weighted sum by $f \approx s_0 + s_1 + s_2 + \cdots$, where the first level s_0 and difference d_1 is defined by

$$s_0(x_k) = \sum_{i=1}^{n} P_0(x_i, x_k) f(x_i), \qquad d_1 = f - s_0,$$

and then iteratively define the lth level decomposition s_l:

$$s_l(x_k) = \sum_{i=1}^{n} P_l(x_i, x_k) d_l(x_i), \qquad d_l = f - \sum_{i=0}^{l-1} s_i.$$

For the choice of kernels k_l, increasing l can lead to overfitting, which we mitigate by zeroing out the diagonals of the kernels (Fernández et al, 2013). We set the truncation level for the iterations at level L once the approximation error begins to increase in l. Given a new data point y, we extend f by

$$\bar{s}_0(y) = \sum_{i=1}^{n} P_0(y, x_i) f(x_i), \qquad \bar{s}_l(y) = \sum_{i=1}^{n} P_l(y, x_i) d_l(x_i),$$

for $L \geq 1$, and assign f the value

$$\bar{f}(y) = \sum_{l=0}^{L} \bar{s}_l(y).$$

That is, we use the kernels P_l to evaluate the similarity of y to points x_i in the training data, and use this measure of similarity to form a weighted average of $f(x_i)$ values to define $\bar{f}(y)$. In practice, not every training data point x_i is used to calculate the weights for this sum, allowing us to ignore the contribution from very dissimilar states that carry very low weight, in addition to reducing computational cost. This restriction is done by zeroing out the smallest entries in $p_y(x)$ and renormalizing by row sum. Note that there will be some reconstruction error between the out-of-sample extension value $\bar{f}(y)$ and the ground truth $f(y)$, which in general is not known.

2.4 Kernel Analog Forecast

Recall that in traditional analog forecasting, a forecast is made by identifying a single historical analog that is most similar to the given initial state, and using the historical analog's trajectory as the forecast. As mentioned in Sect. 2.2, it is convenient to think of normalized kernels as empirical probability distributions in the second argument, $p_y(x) = P(y,x)$. In this framework, traditional analog forecast for a lead time τ can be written as

$$f(y,\tau) = \mathbb{E}_{p_y} S_\tau f = \sum_{i=1}^{n} p_y(x_i) f(x_{i+\tau}) = f(x_{j+\tau}),$$

where $p_y = \delta_{ij}$ and $S_\tau f(x_i) = f(x_{i+\tau})$, the time shifted observable whose value is known on the training dataset.

Given a new state y, we define our prediction for lead time τ, via Laplacian pyramids, by

$$f(y,\tau) = \mathbb{E}_{p_{y,0}} S_\tau f + \sum_{l=1}^{L} \mathbb{E}_{p_{y,l}} S_\tau d_l,$$

where $p_{y,l}(x) = P_l(y,x)$ corresponds to the probability distribution from the kernel at scale l.

The reconstruction error from the out-of-sample extension manifests itself in the fidelity of the forecasts as the error at time lag 0. While in our applications, knowing the climate state y_i allows us to compute the observable exactly $f(y_i)$ at time lag 0, we need the out-of-sample extension to compute the observable $f(y_{i+\tau})$ at any time lag $\tau > 0$, and hence must contend with the initial reconstruction error.

2.5 Error Metrics

We evaluate the performance of our predictions with two aggregate error metrics, the root-mean-square error (RMSE) and pattern correlation (PC), defined as

$$\mathrm{RMSE}^2(\tau) = \frac{1}{n'} \sum_{j=1}^{n'} \left(y_{j+\tau} - x_{j+\tau}\right)^2,$$

$$\mathrm{PC}(\tau) = \frac{1}{n'} \sum_{j=1}^{n'} \frac{\left(y_{j+\tau} - \tilde{y}(\tau)\right)\left(x_{j+\tau} - \tilde{x}(\tau)\right)}{\sigma_y(\tau)\sigma_x(\tau)},$$

where

$$\tilde{y}(\tau) = \frac{1}{n'} \sum_{j=1}^{n'} y_{j+\tau}, \qquad \tilde{x}(\tau) = \frac{1}{n'} \sum_{j=1}^{n'} x_{j+\tau},$$

$$\sigma_y^2(\tau) = \frac{1}{n'} \sum_{j=1}^{n'} (y_{j+\tau} - \tilde{y}(\tau))^2,$$

$$\sigma_x^2(\tau) = \frac{1}{n'} \sum_{j=1}^{n'} (x_{j+\tau} - \tilde{x}(\tau))^2.$$

The averaging is over predictions formed from using testing data of length n' (second half of the data set) as initial conditions. We use the constant persistence forecast $y_\tau = y_0$ as our benchmark, and consider predictive skill to be lost once pattern correlation has dropped below a threshold of 0.5.

3 Datasets

We use CCSM4 (Gent et al, 2011) model data from a pre-industrial control run (b40.1980) and run perfect model prediction experiments, where 800 years of the control run is split into a training dataset and a test dataset, 400 years each. The sea ice component is CICE4 (Hunke and Lipscomb, 2008), the ocean component is POP (Smith et al, 2010), and the atmosphere component is CAM4 (Neale et al, 2010) . Our default experimental setup is to include SIC, SST, and SLP fields, and will later explore the role of SIT as an additional predictor variable. We consider the entire Arctic, as well as the following regions: Beaufort Sea, Chukchi Sea, East Siberian Sea, Laptev Sea, Kara Sea, Barents Sea, Greenland Sea, Baffin Bay, Labrador Sea, Bering Sea, and Sea of Okhotsk. While the ice and ocean state variables are restricted to each region, pan-Arctic SLP data is always used, to allow for possible teleconnection effects. The regions are depicted in Fig. 1, shown with this dataset's sea ice concentration variability.

Since we are using a control run, and predicting anomalies relative to a stationary climatology, all of the variability we see if from natural variability, with no forced variability and no trend. While this is certainly not the case for the current Arctic climate state, our measured forecast skill is absent any assistance of a predicted climatology or trend, and seemingly low skills should be taken in this perspective. An embedding window of 12 months is used; 6 and 24 month embedding windows were also tested for robustness, and while results were similar for a 6 month window, results with 24 months were marginally worse than 12 months.

Our target observable f for prediction is integrated anomalies in sea ice area (as opposed to sea ice extent which is sea ice area above 15% concentration). Sea ice anomalies in the test data period are calculated relative to the monthly climatology calculated from the training data set. While this should not be a concern in a pre-industrial control run with no secular trend, this may be of more importance in other scenarios. Persistence forecasts are initialized with the true anomaly (as opposed to the out-of-sample extension value), so all forecasts will have initial error metrics greater than persistence due to reconstruction error.

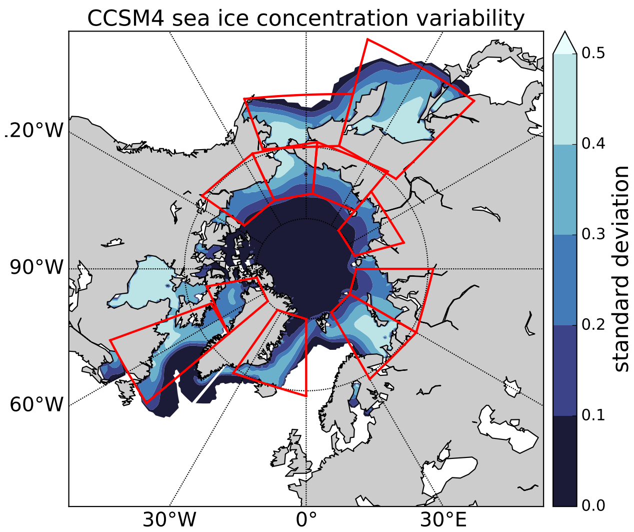

Fig. 1 Standard deviation of sea ice concentration (SIC) for the CCSM4 control run, with regions considered in our forecasting are pan-Arctic (45N–90N), Beaufort Sea (125W–155W, 65N–75N), Chukchi Sea (155W–175E, 65N–75N), East Siberian Sea (175E–140E, 65N–75N), Laptev Sea (140E–105E, 70N–80N), Kara Sea (60E–90E, 65N–80N), Barents Sea (30E–60E, 65N–80N), Greenland Sea (35W–0E, 65N–80N), Baffin Bay (80W–50W, 70N–80N), Labrador Sea (70W–50W, 50–70N), Bering Sea (165E–160W, 55N–65N), and Sea of Okhotsk (135E–165E, 45N–65N).

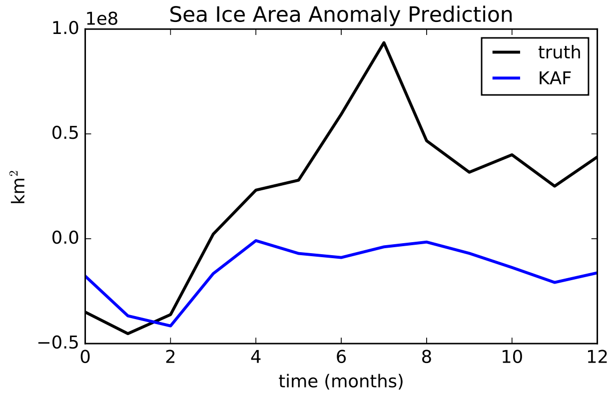

Fig. 2 Sample forecast trajectory of Arctic sea ice cover anomalies. Forecasts typically falter when near a 0 anomaly state, and a considerable amount of error occurs when the forecast moves to the wrong sign when near 0.

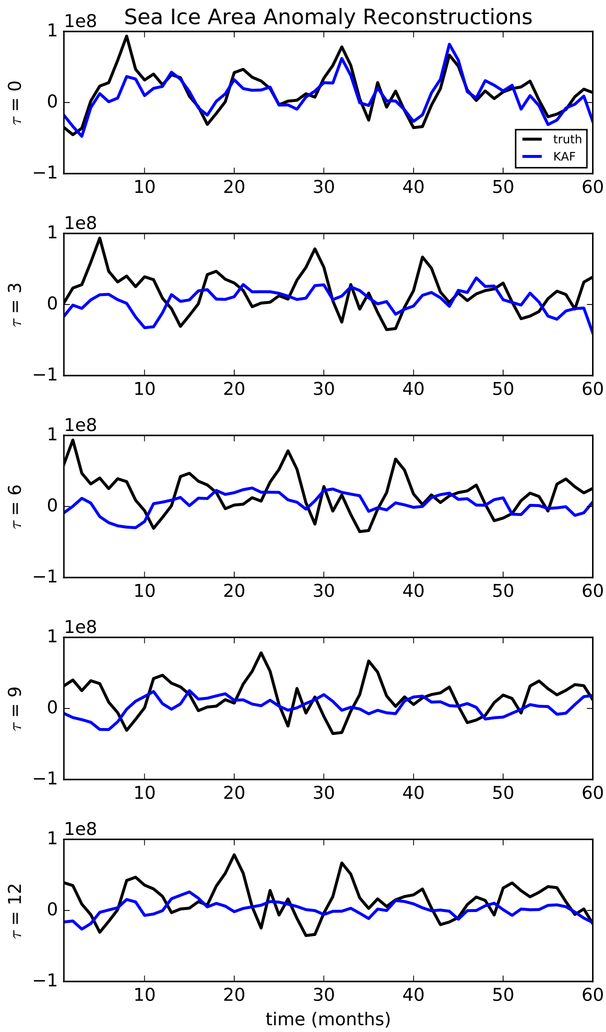

Fig. 3 Sample reconstructions at different time lags for Arctic sea ice area anomalies. Degradation of forecast fidelity is seen as the lead time increases.

4 Results

4.1 Pan-Arctic

We first focus on sea ice area anomalies, using SIC, SST, and SLP initial data, and a 12 month embedding window. Figure 2 shows a sample forecast trajectory compared to the ground truth, initialized in a state relatively strongly away from zero anomaly (climatology). Forecasts typically falter when near zero, as even with dynamic information encoded into the forecasting scheme, there is still difficulty in determining the sign of the forecast anomaly when the state is very near climatology. We show the degradation of the forecasts as lead times increase in Fig. 3, where forecasts are performed with 0, 3, 6, 9, and 12 month lead times. The blue forecast at each point is made from a lead time indicated by the panel. While the initial reconstruction matches the truth reasonably well, forecasts at longer lead times become increasingly smoothed out by averaging, converging to climatology as $\tau \to \infty$.

To quantify how well the forecasts do on average, we consider the error metrics averaged over all points in the test period (400 years of monthly data, minus the length of the embedding window). In Fig. 4, we show pattern correlation conditioned on the target month for prediction and

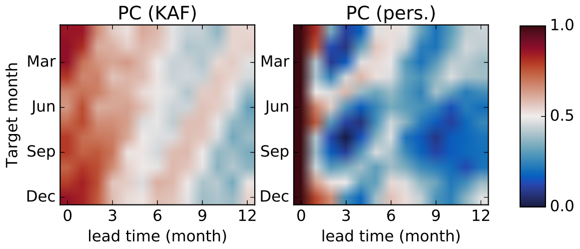

Fig. 4 Pattern correlation as a function of target month of prediction and lead time for pan-Arctic sea ice area anomaly forecasts using KAF (left) and persistence (right). Considering a PC of 0.5 as a threshold for predictive skill, red indicates predictive skill, and blue indicates lack of skill. Beyond the initial reconstruction (lead time 0), KAF outperforms persistence in virtually every other regard. Early summer suffers the lowest predictive skill for short lead times, but is still predictable 4 – 5 months out. For predicting September and March sea ice area anomalies, note that persistence loses skill at short lead times, whereas KAF forecasts show predictive skill 4–6 months out for these months. Note a fall to spring reemergence limb ('growth-to-melt') in the persistence plot.

lead time, for KAF and persistence forecasts as a benchmark for comparison. Beyond initial reconstruction, KAF outperforms persistence in almost every regard, and is even above the 0.5 threshold for almost all of the first 6 months predicted range. Notable is KAF skill in predicting March and September anomalies, successful from 4–6 months out, whereas persistence loses skill in 1–2 months for these local extremes. Late spring is the most troublesome time to predict. We can see evidence of reemergence limbs in the persistence forecast, such as a fall-to-spring limb. Persistence suffers the worst in predicting August/September sea ice anomalies from 3–4 months out, and KAF has noticeably higher scores for predicting September pan-Arctic sea ice area anomaly, which skillful (PC > 0.5) forecasts with lead times out to 9 months.

4.2 Regional Arctic

While predicting pan-Arctic sea ice area minimums and maximums has been of great interest, as more areas of the Arctic become accessible, an increased effort has been made in regional scale predictions. Snapshots of regional ice anomalies (calculated against regional climatologies) in Fig. 5 demonstrate different behavior around the Arctic basin. The out-of-sample extension values are plotted with the truth, and again should be thought of as the lead time 0 forecast. The central Arctic basins (Beaufort, Chukhi, East Siberian, & Laptev, moving clockwise) are largely perennially ice covered, with less variability than other Arctic regions, and we will see that the abundance of time spent near climatology makes predicting anomalies difficult. Continuing clockwise to the Barents Sea, we begin to see the strong influence of the North Atlantic in regulating sea ice cover. More persistent anomalies are seen in the Barents, Greenland, and Labrador seas, which leads to greater predictability. Moving across to the North Pacific basins we have the Bering Sea and Sea of Okhotsk, regions in the marginal ice zone that may spend a couple months of the year completely free of ice. These regions in particular have been found to exhibit strong reemergence phenomena, in both SIC and SST fields (Alexander et al, 1999; Bushuk et al, 2014, 2015; Bushuk and Giannakis, 2015).

The aggregate error metrics, averaged over all months for each region in Figs. 6 (RMSE) and 7 (PC) show that KAF consistently outperforms persistence, (or at least fares no worse) once an initial reconstruction error is overcome, typically after only one month. In pattern correlation, disregarding PC scores below the 0.5 threshold may cut into some apparent gains of KAF over persistence, but it is worth noting the decay rate of KAF PC is slower than persistence, sometimes dramatically so (e.g. Bering, Labrador, Bering). The persistent nature of the North Atlantic adjacent basins seen in Fig. 5 manifests itself as slower than average decay of persistence. Also note the rise in persistence PC for the North Pacific basins (Bering, Okhotsk) at the 12 month mark, suggesting a reemergence phenomena.

Conditioning forecasts on the target month of prediction allows us to parse out seasonal impacts on predictability. The combined spatial and temporal effects of predictability highlight particularly skillful months and regions to predict, as seen in Fig. 8. Regional predictions suffer many lags and initialization months when there is no predictive skill, however times of success are clearly seen. Late summer and fall are in general the most predictable, a pattern that largely appears in each region. Notable is that the September anomaly is the most predictable, but only from a lead time of a couple months. The North Atlantic adjacent regions exhibit larger extent of predictive skill, and some form of reemergence seems to be aiding the forecasts. To demonstrate the gain in predictive skill of KAF over persistence, rather than plot persistence PC by target month as in Fig. 4, we instead plot the difference in pattern correlation scores, KAF over persistence, in Fig. 9. We zero out any value where both PC scores are below the threshold of 0.5, which we consider as not relevant to predictive skill. Most notable improvement over persistence is in the North Atlantic adjacent regions, though many regions also demonstrate improvement in predicting spring anomalies.

The areas of high PC in Fig. 8 by region and season are indicative of the different types of reemergence. In the central Arctic basins (e.g. Beaufort, Chukchi, E. Siberian, Laptev, and Kara), we see regions of high predictability are during the melt seasons, indicating growth-to-melt reemergence is aiding skill. Similarly, in the marginal ice zones

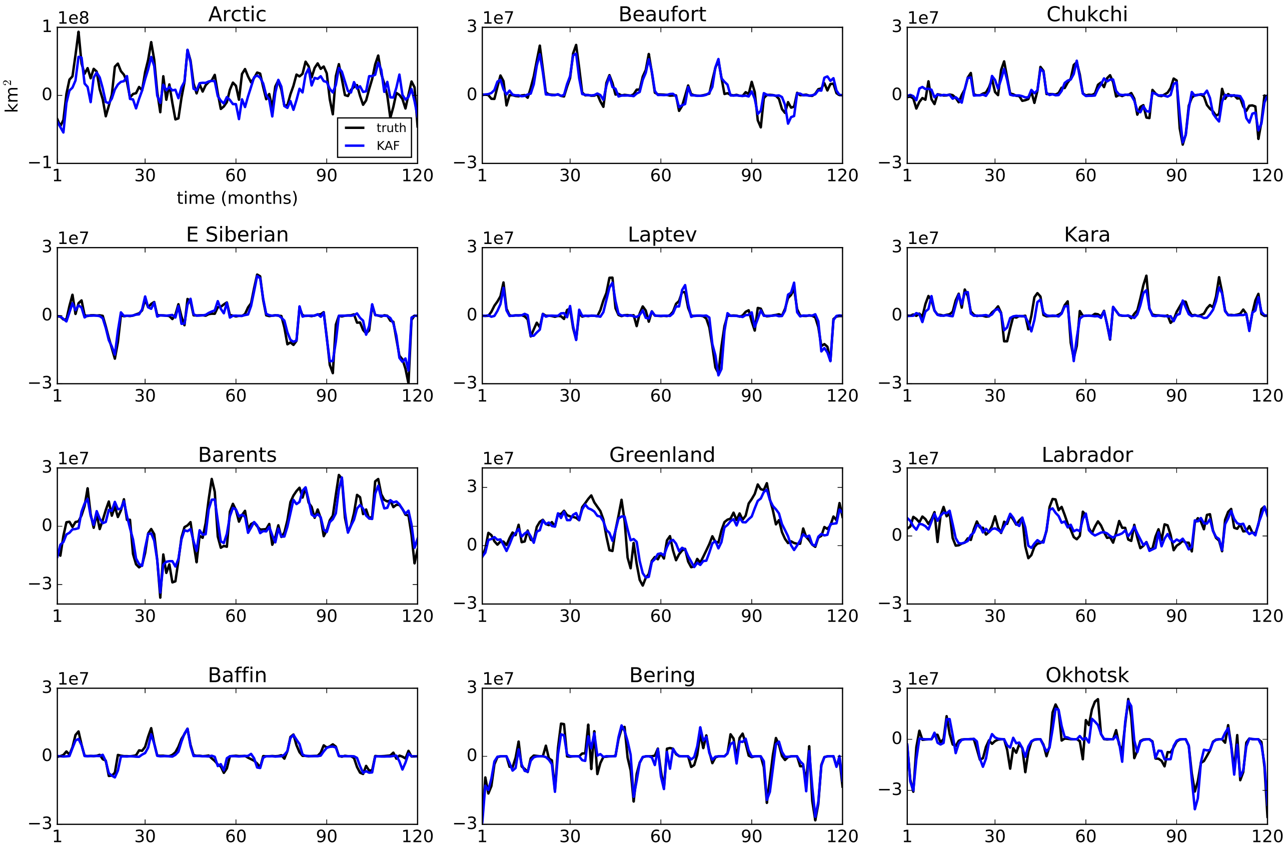

Fig. 5 Initial reconstructions (lead time 0) of regional sea ice area anomalies compared to truth. Note the regions that have more persistent anomalies (Barents, Greenland, and Labrador), are North Atlantic adjacent, which has been found in other studies to be regions of relative high predictability.

of Bering Sea and Sea of Okhotsk, the regions of high predictability are in the growth seasons, where melt-to-growth reemergence is present.

4.3 Role of predictor variables

So far the experiments we have shown have used SIC, SST, and (pan-Arctic) SLP as predictor variables, from which kernel evaluations to determine similarity are based (in Takens embedding space). To address the predictive power of using these components, in Fig. 10 we show the effect of initial data (number of variables used in kernel calculation) on prediction skill, increasingly adding SIC, SST, SLP, and finally SIT. We order the components in this way as roughly increasingly inaccessible as an (near-real time) observational data set, and increasing relevance to the task of sea ice area prediction. While there is marginal difference in adding SST or SLP to SIC, including SIT is actually detrimental to predictions in the pan-Arctic sea ice area anomaly setting. This may seem surprising, given other studies emphasis on the importance of sea ice thickness measurements, however in the context of kernel evaluation, increasing the dimension of our state vector may yield less discernible informative historical analogs. A similar degradation of performance when including SIT data in the kernel was observed in the study of Bushuk and Giannakis (2017) on SIT-SIC reemergence mechanisms. This behavior was attributed to the slower characteristic timescale of SIT data, resulting in their dominating the kernel phase velocity-dependent kernel in Eq. 1.

For regional scale predictions, where SIT may be less variable than in a pan-Arctic setting, it is not the case that incorporating SIT into the state variable impedes predictive skill. In the central Arctic basins, it is typical that thickness adds predictive skill, while the marginal ice zone basins, it is typical that thickness neither helps nor harms predictive skill (Fig. 10).

While the inclusion of pan-Arctic SLP does not hamper our prediction skill, it offers only marginal improvement. This is most likely due to the fact that the quantities used are monthly averaged, and perhaps too temporally coarse to reflect the chaotic atmospheric influence on sea ice cover on shorter time scales.

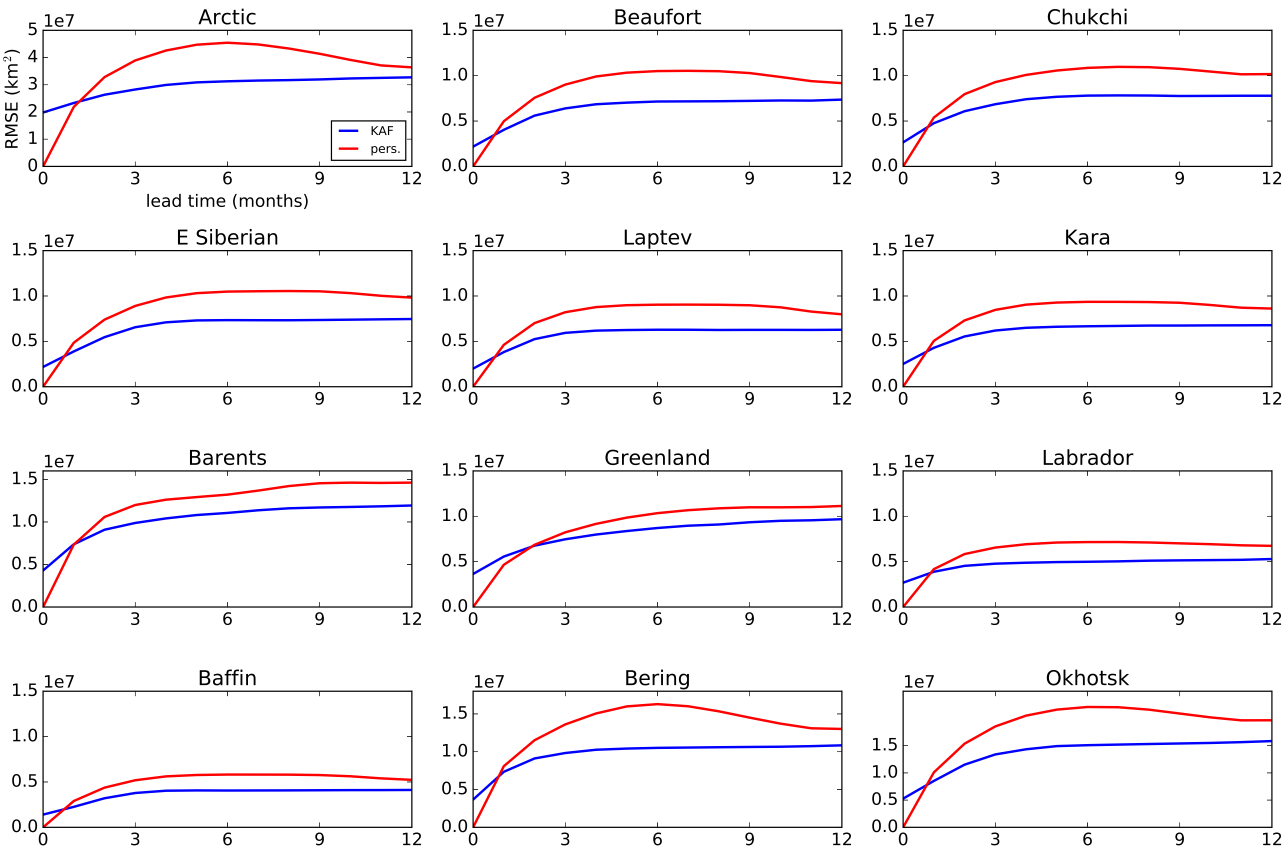

Fig. 6 RMSE for regional sea ice area anomaly predictions, averaged over all months, for KAF and persistence as a benchmark. KAF suffers reconstruction errors at lead time 0, and then outperforms persistence usually after one month.

4.4 Regional volume anomalies

Since we have SIT data available from the CCSM control run, we also consider the problem of forecasting sea ice volume anomalies. In general these show more persistence than sea ice cover anomalies, particularly because thinner ice is more variable and more prone to be advected by winds across region boundaries, or to areas that are more prone to melting. Figure 11 shows regional forecast PC for predicting sea ice volume anomalies, having only observed SIC, SST, and SLP. This is an example of predicting an unobserved variable, and in some areas (late fall, early winter) KAF is remarkably skillful in reconstructing the unobserved quantity, with skill from lead times of a couple months out. However these do not compare favorably against a persistence forecast using the ground truth (not shown), due to inherent persistence of volume anomalies (compared to area anomalies which persist on a shorter timescale), though this would also not be a fair comparison given the KAF forecasts are not observing the full observable. We note that in comparing to Fig. 8 we see similar patterns of predictive skill by region, suggesting that again the different types of reemergence are aiding in skill. When SIT is included in the initial data, as in Fig. 12, the regional predictive skill is extended to 3–6 months, with particular improvement in the central Arctic basin regions. Again, however, persistence is not exceeded in forecast skill.

5 Discussion & Conclusions

In this paper, we used KAF (Zhao and Giannakis, 2016; Comeau et al, 2017; Alexander et al, 2017), a nonparametric method using weighted ensembles of analogs, to predict Arctic sea ice area anomalies, then volume anomalies, for both pan-Arctic and regional scales, examining the effects of including SIC, SST, SLP, and SIT as predictors for our method. We find in general that KAF outperforms the persistence forecast, or at minimum does not perform worse, with outperformance lead times ranging between 0 and 9 months. Moving to regional scale basins and conditioning on the target month of prediction, we see clear regional-seasonal domains when KAF succeeds, as well as those when it fails (along with persistence).

The North Atlantic seems to have a strong impact on sea ice area anomalies, as the adjacent regions (Barents & Kara,

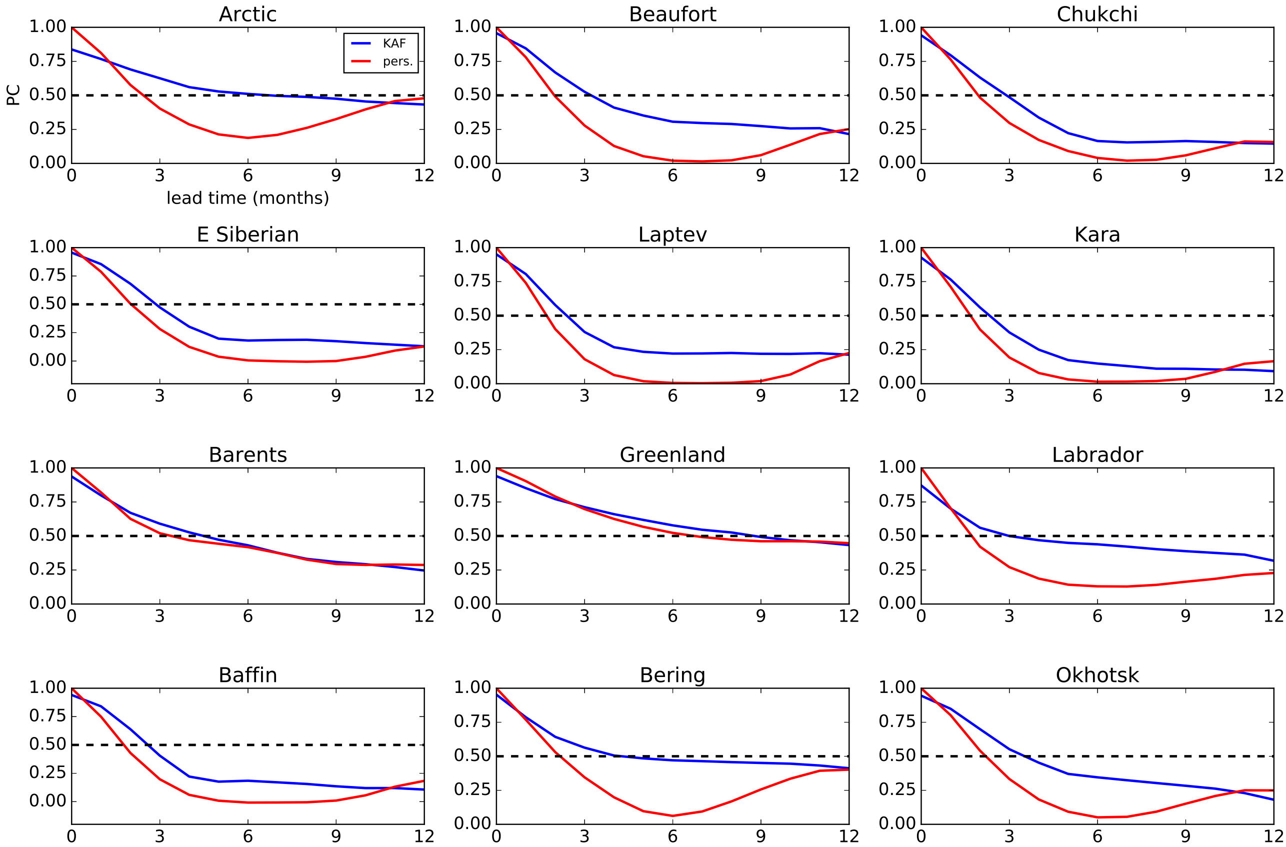

Fig. 7 Pattern correlation scores for regional sea ice area anomaly prediction by region, averaged over all months. Reconstruction errors are less noticeable in this metric (apart from pan-Arctic), and KAF exceeds persistence in every region. Note the North Atlantic regions are the most persistent, as demonstrated by slow decay of PC. Regional improvements over persistence are marginal, though in pan-Arctic, the improvement is several months.

Greenland, and Labrador Seas) exhibit the strongest persistent anomalies, and have the highest predictability. The marginal ice zone basins in the North Pacific (Bering Sea, Sea of Okhotsk) show similar behavior, but to a lesser degree. The basins in the central Arctic (Beaufort, Chukchi, East Siberian, and Laptev) have less variability to their sea ice cover, and thus being close to climatology for much of the time makes predicting excursions from the climatology difficult. Late summer and early fall are in general the best times to predict with KAF, with skillful forecasts at lead times 3–6 months. Late winter and spring is in general the time period of least predictability, except in the marginal ice zones. The areas of high predictability by region and season are consistent with reemergence phenomena, with central Arctic basins benefiting from melt to growth reemergence, and marginal ice zones benefiting from growth to melt reemergence.

We find most of the predictive information is in SIC and SST, with possible marginal improvements in incorporating SLP and SIT. The marginal improvements in these predictor variables is possibly due to coarse temporal resolution in the atmosphere component, or in approaching too high of dimension where the kernel begins to fail to identify useful historical analogs. While we have success in reconstructing sea ice volume anomalies without observing thickness, this cannot compete with true knowledge of the system in a forecast setting, given the persistent nature of volume anomalies.

Ultimately, the goal is to move to an operational prediction based on observational data, for which this is a first step. While this method could be applied to the observational data itself (without a removed climatology), we would not expect very good results given the quite reasonable caution that others have made on basing a statistical prediction on historical relationships in a changing mean Arctic state, which will almost certainly overestimate future Arctic sea ice area. Our future research plan is to use a nonlinear dimension reduction method to extract an underlying 'trend' in the data as a way of non-parametrically determining a trend (as opposed to fitting a linear or quadratic regression). This trend could then be extended to a forecast time using some form of extrapolation or out-of-sample extension technique, while the anomalies from this trend would be forecasted by the KAF method discussed in this study. Other research directions include using a blended persistence and analog forecasting ap-

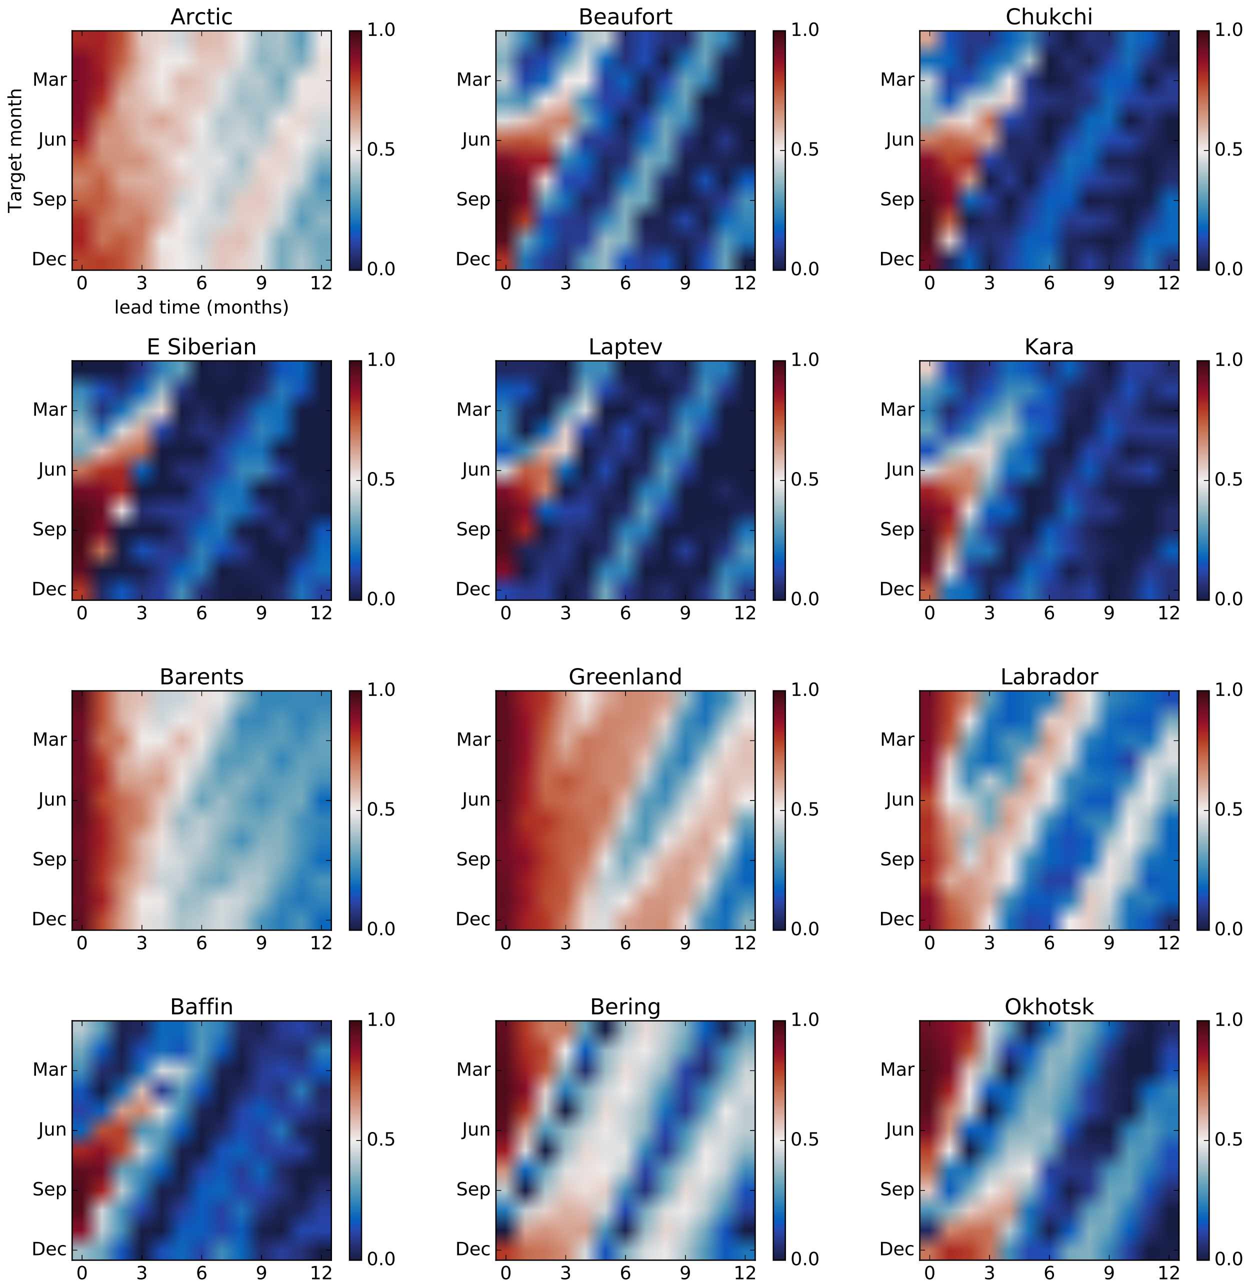

Fig. 8 Pattern correlation as a function of target month and lead time by region. Notice pan-Arctic Sep seems lightest color for lag 0, but is over 0.5 for 3 – 4 months. In general across regions, predicting fall is typically the most skillful, with the exception of the North Pacific basins in marginal ice zones (Bering Sea and Sea of Okhotsk), where higher predictability is in the winter. The North Atlantic adjacent regions have a larger extent of predictive skill, with apparent limbs of high correlation that may be reemergence processes aiding KAF.

proach to avoid the initial reconstruction errors at short time scales, as well as forecasts using kernels targeted at specific observables.

Acknowledgements The research of Andrew Majda and Dimitrios Giannakis is partially supported by ONR MURI grant 25-74200-F7112. Darin Comeau was supported as a postdoctoral fellow through this grant. Dimitrios Giannakis and Zhizhen Zhao are partially supported by NSF grant DMS-1521775. Dimitrios Giannakis also acknowledges support from ONR grant N00014-14-1-0150. We thank Mitch Bushuk for helpful discussions.

References

Alessandrini S, Delle Monache L, Sperati S, Nissen J (2015) A novel application of an analog ensemble for short-term wind power forecasting. Renewable Energy 76:768–781

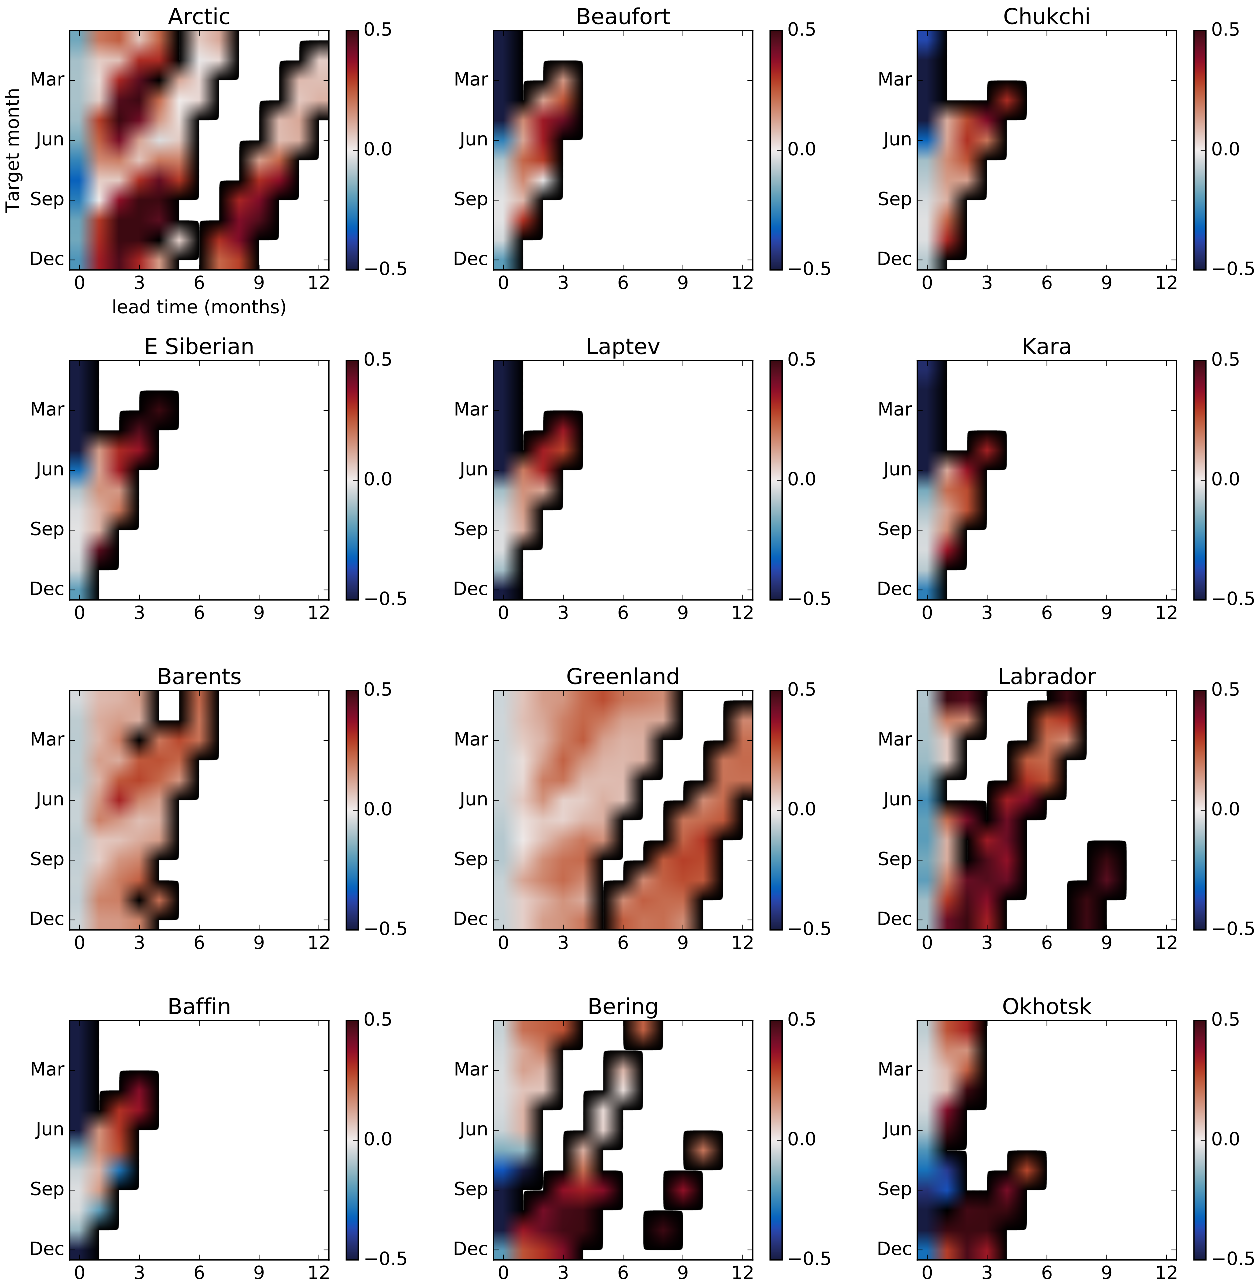

Fig. 9 Difference in pattern correlation scores of KAF over persistence to illustrate the gain in predictive skill, with zero in place of any value where both scores are below 0.5. Red indicates KAF outperforming persistence, and blue vice-versa. For pan-Arctic, a big improvement over persistence is observed late in the year, with a gap representing a reemergence phenomena. Strong improvement is observed in predictability in North Atlantic adjacent regions, as well as in predicting (early) spring anomalies in many regions.

Alexander MA, Deser C, Timlin MS (1999) The reemergence of SST anomalies in the North Pacific Ocean. Journal of Climate 12(8):2419–2433

Alexander R, Zhao Z, Szkely E, Giannakis D (2017) Kernel analog forecasting of tropical intraseasonal oscillations. Journal of the Atmospheric Sciences 74(4):1321–1342

Blanchard-Wrigglesworth E, Bitz CM (2014) Characteristics of Arctic sea-ice thickness variability in GCMs. Journal of Climate 27(21):8244–8258

Blanchard-Wrigglesworth E, Armour KC, Bitz CM, DeWeaver E (2011a) Persistence and inherent predictability of Arctic sea ice in a GCM ensemble and observations. Journal of Climate 24(1):231–250

Blanchard-Wrigglesworth E, Bitz C, Holland M (2011b) Influence of initial conditions and climate forcing on predicting Arctic sea ice. Geophysical Research Letters

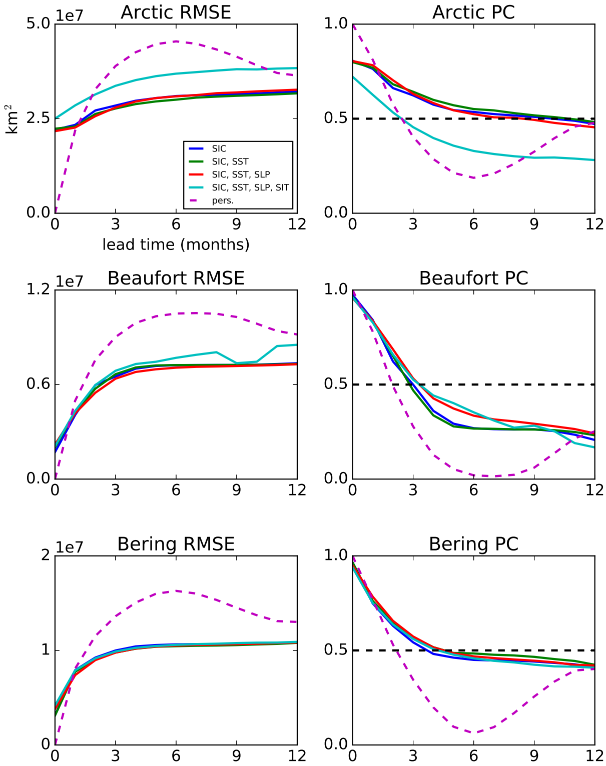

Fig. 10 Prediction results for Arctic sea ice area anomalies using different predictor variables for the pan-Arctic, a region in a mainly perennial ice zone (Beaufort Sea), and a region in a marginal ice zone (Bering Sea). Most of the skill is from SIC alone, although SIC and SST performs (marginally) the best. Interestingly, adding ice thickness information actually hinders pan-Arctic sea ice area anomaly prediction.

38(18)

Blanchard-Wrigglesworth E, Cullather R, Wang W, Zhang J, Bitz C (2015) Model forecast skill and sensitivity to initial conditions in the seasonal sea ice outlook. Geophysical Research Letters 42(19):8042–8048

Broomhead DS, King GP (1986) Extracting qualitative dynamics from experimental data. Physica D: Nonlinear Phenomena 20(2-3):217–236

Bushuk M, Giannakis D (2015) Sea-ice reemergence in a model hierarchy. Geophysical Research Letters 42(13):5337–5345

Bushuk M, Giannakis D (2017) The seasonality and interannual variability of arctic sea ice reemergence. Journal of Climate (2017)

Bushuk M, Giannakis D, Majda AJ (2014) Reemergence mechanisms for North Pacific sea ice revealed through nonlinear Laplacian spectral analysis*. Journal of Climate 27(16):6265–6287

Bushuk M, Giannakis D, Majda AJ (2015) Arctic sea ice reemergence: The role of large-scale oceanic and atmospheric variability. Journal of Climate 28(14):5477–5509

Bushuk M, Msadek R, Winton M, Vecchi GA, Gudgel R, Rosati A, Yang X (2017) Summer enhancement of arctic sea ice volume anomalies in the september-ice zone. Journal of Climate 30(7):2341–2362

Chevallier M, Salas-Mélia D (2012) The role of sea ice thickness distribution in the Arctic sea ice potential predictability: A diagnostic approach with a coupled gcm. Journal of Climate 25(8):3025–3038

Chevallier M, Salas y Mélia D, Voldoire A, Déqué M, Garric G (2013) Seasonal forecasts of the pan-Arctic sea ice extent using a GCM-based seasonal prediction system. Journal of Climate 26(16):6092–6104

Comeau D, Zhao Z, Giannakis D, Majda AJ (2017) Data-driven prediction strategies for low-frequency patterns of North Pacific climate variability. Climate Dynamics 48(5):1855–1872

Day J, Tietsche S, Hawkins E (2014) Pan-Arctic and regional sea ice predictability: Initialization month dependence. Journal of Climate 27(12):4371–4390

Deyle ER, Sugihara G (2011) Generalized theorems for nonlinear state space reconstruction. PLoS One 6(3):e18,295

Drosdowsky W (1994) Analog (nonlinear) forecasts of the Southern Oscillation index time series. Weather and Forecasting 9(1):78–84

Fernández A, Rabin N, Fishelov D, Dorronsoro JR (2013) Auto-adaptative laplacian pyramids for high-dimensional data analysis. arXiv preprint arXiv:13116594

Gent PR, Danabasoglu G, Donner LJ, Holland MM, Hunke EC, Jayne SR, Lawrence DM, Neale RB, Rasch PJ, Vertenstein M, et al (2011) The community climate system model version 4. Journal of Climate 24(19):4973–4991

Giannakis D (2015) Dynamics-adapted cone kernels. SIAM Journal on Applied Dynamical Systems 14(2):556–608

Giannakis D, Majda AJ (2012a) Comparing low-frequency and intermittent variability in comprehensive climate models through nonlinear Laplacian spectral analysis. Geophysical Research Letters 39(10)

Giannakis D, Majda AJ (2012b) Nonlinear Laplacian spectral analysis for time series with intermittency and low-frequency variability. Proceedings of the National Academy of Sciences 109(7):2222–2227

Giannakis D, Majda AJ (2013) Nonlinear Laplacian spectral analysis: capturing intermittent and low-frequency spatiotemporal patterns in high-dimensional data. Statistical Analysis and Data Mining 6(3):180–194

Giannakis D, Majda AJ (2014) Data-driven methods for dynamical systems: Quantifying predictability and extracting spatiotemporal patterns. Mathematical and Computational Modeling: With Applications in Engineering and the Natural and Social Sciences p 288

Guemas V, Blanchard-Wrigglesworth E, Chevallier M, Day JJ, Déqué M, Doblas-Reyes FJ, Fučkar NS, Germe A,

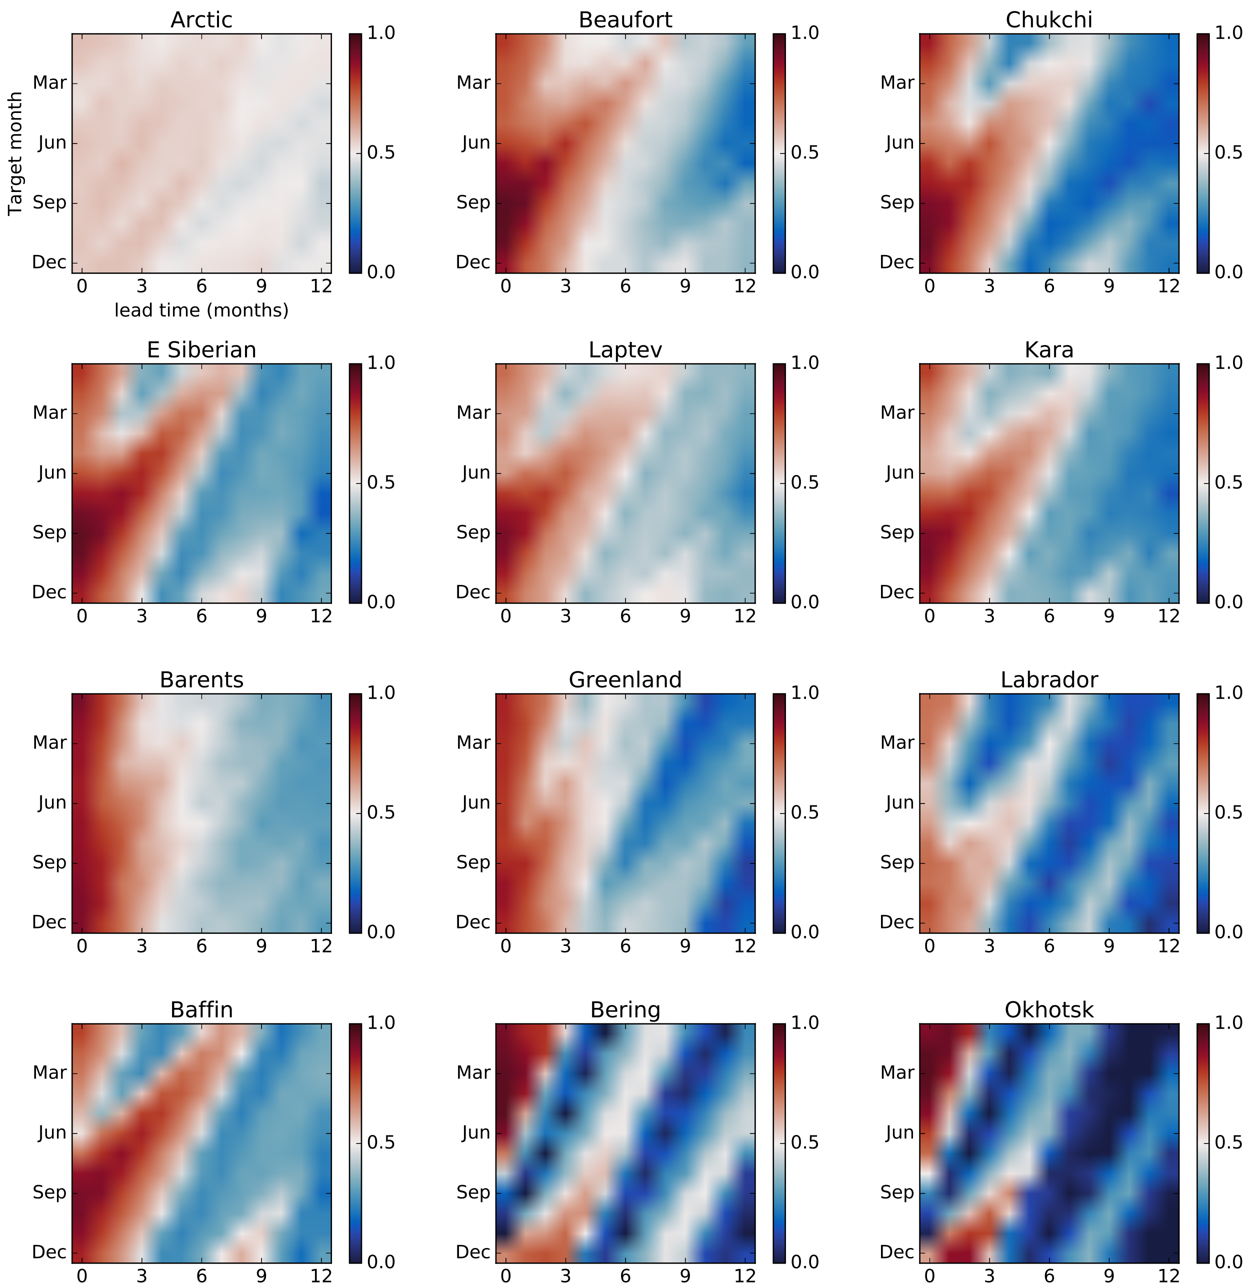

Fig. 11 Forecasts for Arctic sea ice volume anomalies, predicted using only SIC, SST, and SLP data. By not using SIT as a predictor variable, we are predicting an unobserved quantity. Given the longer timescales of sea ice volume anomalies (due to thermodynamics playing a larger role), persistence forecasts based on the true anomaly outperform KAF on all spatial and temporal scales (not shown).

Hawkins E, Keeley S, et al (2014) A review on arctic sea-ice predictability and prediction on seasonal to decadal time-scales. Quarterly Journal of the Royal Meteorological Society

Holland MM, Stroeve J (2011) Changing seasonal sea ice predictor relationships in a changing Arctic climate. Geophysical Research Letters 38(18)

Holland MM, Bitz CM, Tremblay L, Bailey DA, et al (2008) The role of natural versus forced change in future rapid summer Arctic ice loss. Arctic sea ice decline: observations, projections, mechanisms, and implications pp 133–150

Holland MM, Bailey DA, Vavrus S (2011) Inherent sea ice predictability in the rapidly changing Arctic environment of the Community Climate System Model, version 3. Climate Dynamics 36(7-8):1239–1253

Holland MM, Bailey DA, Briegleb BP, Light B, Hunke E (2012) Improved sea ice shortwave radiation physics in

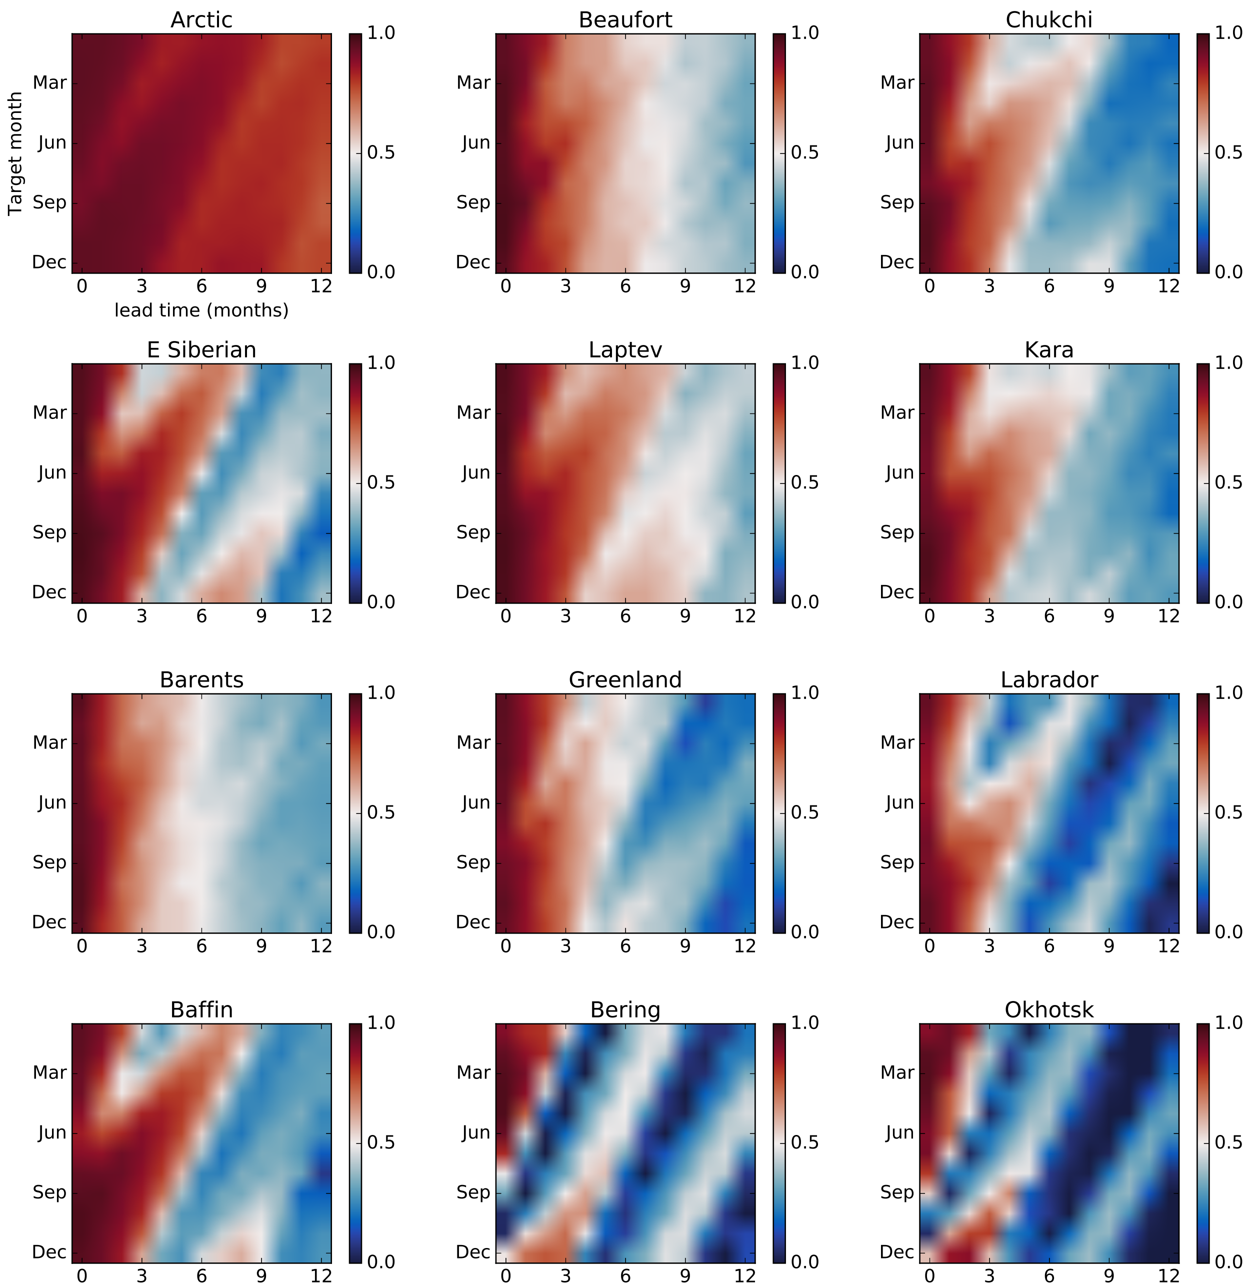

Fig. 12 Same as Fig. 11, with the inclusion of SIT as a predictor variable, granting a boost to forecast skill, particularly in the central Arctic basin regions.

CCSM4: the impact of melt ponds and aerosols on Arctic sea ice*. Journal of Climate 25(5):1413–1430

Hunke E, Lipscomb W (2008) CICE: The Los Alamos sea ice model, documentation and software, version 4.0, Los Alamos National Laboratory tech. rep. Tech. rep., LA-CC-06-012

Koenigk T, Mikolajewicz U (2009) Seasonal to interannual climate predictability in mid and high northern latitudes in a global coupled model. Climate dynamics 32(6):783–798

Lindsay R, Zhang J, Schweiger A, Steele M (2008) Seasonal predictions of ice extent in the Arctic Ocean. Journal of Geophysical Research: Oceans 113(C2)

Lorenz EN (1969) Atmospheric predictability as revealed by naturally occurring analogues. Journal of the Atmospheric sciences 26(4):636–646

Neale RB, Richter JH, Conley AJ, Park S, Lauritzen PH, Gettelman A, Williamson DL, et al (2010) Description of the ncar community atmosphere model (cam 4.0). NCAR Tech Note NCAR/TN-485+ STR

Ogi M, Yamazaki K, Wallace JM (2010) Influence of winter and summer surface wind anomalies on summer Arctic sea ice extent. Geophysical Research Letters 37(7)

Packard NH, Crutchfield JP, Farmer JD, Shaw RS (1980) Geometry from a time series. Physical Review Letters 45(9):712

Rabin N, Coifman RR (2012) Heterogeneous datasets representation and learning using diffusion maps and Laplacian pyramids. In: SDM, SIAM, pp 189–199

Robinson JC (2005) A topological delay embedding theorem for infinite-dimensional dynamical systems. Nonlinearity 18(5):2135

Sauer T, Yorke JA, Casdagli M (1991) Embedology. Journal of statistical Physics 65(3-4):579–616

Schröder D, Feltham DL, Flocco D, Tsamados M (2014) September Arctic sea-ice minimum predicted by spring melt-pond fraction. Nature Climate Change 4(5):353–357

Sigmond M, Reader MC, Flato GM, Merryfield WJ, Tivy A (????) Skillful seasonal forecasts of arctic sea ice retreat and advance dates in a dynamical forecast system. Geophysical Research Letters 43(24)

Sigmond M, Fyfe J, Flato G, Kharin V, Merryfield W (2013) Seasonal forecast skill of Arctic sea ice area in a dynamical forecast system. Geophysical Research Letters 40(3):529–534

Smith LC, Stephenson SR (2013) New trans-Arctic shipping routes navigable by midcentury. Proceedings of the National Academy of Sciences 110(13):E1191–E1195

Smith R, Jones P, Briegleb B, Bryan F, Danabasoglu G, Dennis J, Dukowicz J, Eden C, Fox-Kemper B, Gent P, et al (2010) The Parallel Ocean Program (POP) reference manual: Ocean component of the Community Climate System Model (CCSM). Los Alamos National Laboratory, LAUR-10-01853

Stroeve J, Hamilton LC, Bitz CM, Blanchard-Wrigglesworth E (2014) Predicting September sea ice: Ensemble skill of the SEARCH sea ice outlook 2008–2013. Geophysical Research Letters 41(7):2411–2418

Stroeve JC, Kattsov V, Barrett A, Serreze M, Pavlova T, Holland M, Meier WN (2012) Trends in Arctic sea ice extent from CMIP5, CMIP3 and observations. Geophysical Research Letters 39(16)

Takens F (1981) Detecting strange attractors in turbulence. Springer

Tietsche S, Notz D, Jungclaus JH, Marotzke J (2013) Predictability of large interannual Arctic sea-ice anomalies. Climate dynamics 41(9-10):2511–2526

Tietsche S, Day J, Guemas V, Hurlin W, Keeley S, Matei D, Msadek R, Collins M, Hawkins E (2014) Seasonal to interannual Arctic sea ice predictability in current global climate models. Geophysical Research Letters 41(3):1035–1043

Wang L, Yuan X, Ting M, Li C (2015) Predicting summer Arctic sea ice concentration intra-seasonal variability using a vector auto-regressive model. Journal of Climate (2015)

Wang W, Chen M, Kumar A (2013) Seasonal prediction of Arctic sea ice extent from a coupled dynamical forecast system. Monthly Weather Review 141(4):1375–1394

Xavier PK, Goswami BN (2007) An analog method for real-time forecasting of summer monsoon subseasonal variability. Monthly Weather Review 135(12):4149–4160

Zhang J, Rothrock D (2003) Modeling global sea ice with a thickness and enthalpy distribution model in generalized curvilinear coordinates. Monthly Weather Review 131(5):845–861

Zhang J, Steele M, Lindsay R, Schweiger A, Morison J (2008) Ensemble 1-year predictions of Arctic sea ice for the spring and summer of 2008. Geophysical Research Letters 35(8)

Zhao Z, Giannakis D (2016) Analog forecasting with dynamics-adapted kernels. Nonlinearity 29(9):2888