Monthly Sea-Surface Temperature, Sea Ice, and Sea-Level Pressure over 1850–2023 from Coupled Data Assimilation

Vincent T. Cooper,[a] Gregory J. Hakim,[a] Kyle C. Armour,[a,b]

[a] *Department of Atmospheric and Climate Science, University of Washington, Seattle, WA, USA*
[b] *School of Oceanography, University of Washington, Seattle, WA, USA*

ABSTRACT: Historical observations of Earth's climate underpin our knowledge and predictions of climate variability and change. However, the observations are incomplete and uncertain, and existing reanalysis datasets based on these observations are derived separately for each component of the climate system, yielding inconsistencies that limit understanding of coupled climate dynamics. Here we use coupled data assimilation, which synthesizes observational and dynamical constraints across all climate fields simultaneously, to reconstruct globally resolved sea-surface temperature (SST), near-surface air temperature (T), sea-level pressure (SLP), and sea-ice concentration (SIC), over 1850–2023. We use a Kalman filter and forecasts from an efficient emulator (Linear Inverse Model; LIM) to assimilate observations of SST, land T, marine SLP, and satellite-era SIC. We account for model error by training LIMs on eight CMIP6 models, and we use the LIMs to generate eight independent reanalyses with 200 ensemble members, yielding 1600 total members. Key findings in the Tropics include post-1980 trends in the Walker circulation that are consistent with past variability, whereas the tropical SST contrast (the difference between warmer and colder SSTs) shows a distinct strengthening since 1975. ENSO amplitude exhibits substantial low-frequency variability and a local maximum in variance from 1875–1910. In polar regions, we find a muted cooling trend in the Southern Ocean post-1980 and substantial uncertainty. Changes in Antarctic sea ice are relatively small between 1850 and 2000, while Arctic sea ice declines by 0.5 ± 0.1 (1σ) million km^2 during the 1920s.

SIGNIFICANCE STATEMENT: The key advance in our reconstruction is that the ocean, atmosphere, and sea ice are dynamically consistent with each other and with observations across all components, thus forming a true climate reanalysis. Existing climate datasets are derived separately for each component (e.g., atmosphere, ocean, and sea ice), leading to spurious trends and inconsistencies in coupled climate variability. We use coupled data assimilation to unify observations and coupled dynamics across components. We combine forecasts from climate models with observations from ocean vessels and weather stations to produce monthly state estimates spanning 1850–2023 and a novel quantification of globally resolved uncertainty. This reconstruction provides insights into historical variability and trends while motivating future efforts to reduce uncertainties in the climate record.

1. Introduction

The historical record (c. 1850–present) is central to our understanding of climate variability and Earth's response to anthropogenic forcings, but we have yet to fully extract the available information from instrumental data. Observations of sea-surface temperature (SST), near-surface air temperature (T), and sea-level pressure (SLP) from ships of opportunity and weather stations are noisy, sparse, and vary over time, which adds an incomplete-data problem (Schneider 2001) to analyses of climate variability and change that cannot be avoided and should not be ignored.

To prepare observations for climate analysis, data sources must first be homogenized (e.g., Kent and Kennedy 2021; Chan and Huybers 2019; Chan et al. 2023; Karl et al. 2015; Hausfather et al. 2017; Cowtan et al. 2018), and then the missing values must be imputed. Imputation is typically statistical, employing pattern-based methods (including empirical orthogonal functions; EOFs) derived from recent decades or kriging, and does not involve dynamical constraints (e.g., Kaplan et al. 1998; Rohde et al. 2013; Cowtan and Way 2014; Hirahara et al. 2014; Huang et al. 2017; Kadow et al. 2020; Vaccaro et al. 2021). Furthermore, when values are imputed for different climate fields, e.g., SST and SLP, there are no dynamical constraints ensuring that the coupled fields are physically consistent. Imputation and homogenization can have pronounced impacts on assessments of the climate sensitivity to increasing greenhouse gases (e.g., Sherwood et al. 2020; Forster et al. 2021; Modak and Mauritsen 2023), efforts to distinguish internal variability from forced climate change (e.g., Hegerl et al. 2019; Wills et al. 2020), understanding of atmosphere–ocean variability (e.g., Battisti et al. 2019), and evaluation of climate models (e.g., Wills et al. 2022). Here we apply a different approach to solve the incomplete-data problem: we use coupled data assimilation to impose observational and dynamical constraints across all climate fields simultaneously, ensuring that the full state estimate is internally consistent.

SST patterns play a ubiquitous role in regulating climate variability (e.g., Bjerknes 1969; Barsugli and Battisti 1998; Alexander et al. 2002; Deser et al. 2010a; Newman et al.

Corresponding author: Vincent T. Cooper, vcooper@uw.edu

2016; Czaja et al. 2019; Capotondi et al. 2023), radiative feedbacks (e.g., Armour et al. 2013; Andrews et al. 2015; Zhou et al. 2016; Ceppi and Gregory 2017; Andrews and Webb 2018; Dong et al. 2019; Fueglistaler 2019; Andrews et al. 2022; Salvi et al. 2023; Cooper et al. 2024), and the hydrologic cycle (Hastenrath and Greischar 1993; Xie et al. 2010; Hoerling et al. 2010; Chadwick et al. 2014; Lehner et al. 2018; Siler et al. 2019; Cook et al. 2022; Kuo et al. 2023; Seager et al. 2023). There are a variety of recent (c. 1980–present) climate phenomena tied to SSTs that seem either unprecedented or unremarkable depending on what we deem to be natural variability, and this interpretation of recent trends relies on the incomplete and brief instrumental record (e.g., Wunsch 1999). In the tropical Pacific, the zonal SST gradient has strengthened (Solomon and Newman 2012; Coats and Karnauskas 2017; Lee et al. 2022; Watanabe et al. 2024), with cooling in the East Pacific and warming in the West Pacific that has coincided with an apparent strengthening of the Walker circulation post-1980 (e.g., L'Heureux et al. 2013; McGregor et al. 2014; Watanabe et al. 2023; Heede and Fedorov 2023) and distinct tropospheric temperature trends (Flannaghan et al. 2014; Fueglistaler 2019; Po-Chedley et al. 2021). At the poles, the Arctic has warmed rapidly since 1980 with substantial loss of sea ice (Dörr et al. 2023; England et al. 2021; Notz and SIMIP Community 2020), while the Southern Ocean has cooled with an overall expansion of sea ice—until 2015, after which the Southern Ocean has shown rapid warming and sea ice loss (Fan et al. 2014; Stuecker et al. 2017; Fogt et al. 2022; Espinosa et al. 2024; Roach and Meier 2024; Zhang and Li 2023; Turner et al. 2022; Dong et al. 2023; Suryawanshi et al. 2023; Bonan et al. 2024).

A major challenge for the climate dynamics community is understanding the causes of these observed changes as well as the apparent yet debated inability of our state-of-the-art coupled climate models to replicate them (e.g., Wills et al. 2022; Dong et al. 2021; Rugenstein et al. 2023; Seager et al. 2022; Olonscheck et al. 2020; Chung et al. 2019; Watanabe et al. 2021; Roach et al. 2020; Notz and SIMIP Community 2020; Chemke et al. 2022; Kang et al. 2024). Progress on this endeavor requires robustly quantifying observational uncertainties and placing recent changes in historical context with more reliable reconstructions of past climate changes. For example, are the post-1980 trends in tropical SST gradients, the Walker circulation, and polar climates unique over the historical record, or have such changes occurred often due to internal climate variability?

Existing SST datasets designed for climate analysis use a variety of statistical interpolation methods. These methods have been recently summarized in Modak and Mauritsen (2023) and Lewis and Mauritsen (2021) and described in detail in a review by Kent and Kennedy (2021), which also explains the extensive efforts to homogenize time-varying sources of in situ data. To assess the atmospheric response to SST changes over the historical record in atmospheric general circulation models (i.e., in AMIP-type simulations; Eyring et al. 2016; Webb et al. 2017), complete coverage and monthly resolution of SST and sea-ice concentration (SIC) is required. Combined SST and SIC datasets for this purpose include the 1870–2022 PCMDI AMIP II boundary condition (Hurrell et al. 2008) used as the standard for CMIP6, 1854–present NOAA ERSSTv5 (Huang et al. 2017), Met Office Hadley Centre's 1870–present HadISST1 (Rayner et al. 2003) and 1850–2010 HadISST2.1 (no longer maintained; Titchner and Rayner 2014), and the Japanese Meteorological Agency's 1850–present COBE-SST2 (Hirahara et al. 2014). Kaplan et al. (1998) developed a landmark SST analysis using optimal interpolation, and since then the incomplete-data problem has been addressed using kriging (Cowtan and Way 2014), Markov random graphs (Vaccaro et al. 2021), and machine learning (Kadow et al. 2020) to impute hybrid air-sea surface temperatures over land and ocean.

Figure 1 depicts the time-evolving observing network of in situ SST measurements in HadSST4 (Kennedy et al. 2019). As motivation for this study, we illustrate the spread (1σ) across existing datasets (HadISST1, HadISST2.1, ERSSTv5, COBE-SST2, and AMIPII) in their preindustrial-baseline SST (mean anomaly over years 1870–1899) and the spread in their SST trends from 1900–1979 and 1980–2010. We separate the satellite era (c. 1980–present) from the earlier warming because of the variety of studies highlighting and questioning the peculiarity of recent SST trends (e.g., Fueglistaler and Silvers 2021; Andrews et al. 2022; Lewis and Mauritsen 2021). The spatial pattern of uncertainty is influenced by varying methods of imputation, homogenization of data sources, and representativeness error in using point observations as estimates of grid-scale means. Even after 1980, the data coverage over the Southern Ocean and southeast Pacific is notably far from complete, and the inter-dataset differences in those regions are substantial in recent decades (Figure 1c,f).

Atmospheric reanalyses address the incomplete-data problem with data assimilation, which uses a weather model's dynamics to constrain the atmospheric state. Data assimilation (DA) broadly describes the collection of methods that synthesize model forecasts with sparse and noisy observations, producing posterior analyses and uncertainties that are subject to the dynamical constraints of the model. DA is computationally intensive, hence existing reanalyses only assimilate atmospheric observations and only apply dynamical constraints to the atmospheric component, meaning that the SST and SIC boundary conditions are prescribed a priori in ERA5 (Hersbach et al. 2020), JRA-55 (Kobayashi et al. 2015), NOAA's 20$^{\text{th}}$ Century Reanalysis (Compo et al. 2011; Slivinski et al. 2019), and

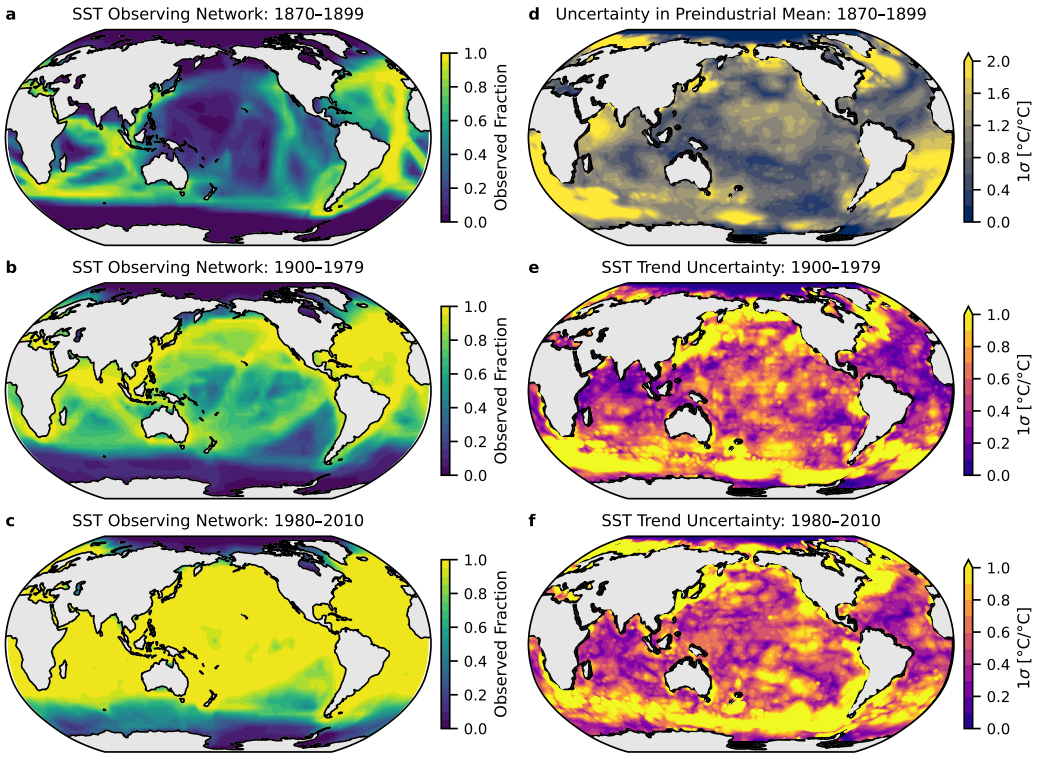

FIG. 1. **Historical observing network and SST uncertainty in pre-existing infilled datasets. (a–c)** Fraction of months with in situ data for SST over three time periods in HadSST4, where 1.0 indicates data in every month during the period. **(d)** Illustration of systematic uncertainty in normalized pattern of preindustrial-mean SST anomalies across existing infilled datasets, calculated as the sample standard deviation (1σ) of the 1870–1899 mean anomalies across HadISST1, HadISST2.1, ERSSTv5, PCMDI/AMIP II, and COBE-SST2, relative to their 1961–1990 climatologies; local anomalies are divided by global-mean anomalies (60°S–60°N) to highlight uncertainty in spatial patterns. **(e–f)** Illustration of systematic uncertainty in patterns of SST trends, calculated as the 1σ of local trends across the same datasets in panel **d**; local SST trends are first divided by the global-mean SST trends (60°S–60°N) to highlight uncertainty in the patterns, and local values greater than 1.0 indicate that the local 1σ is greater than the global-mean trend. Note different colorbars in panels **d–f**.

Mod-ERA (Franke et al. 2017; Valler et al. 2024). Coupled atmosphere–ocean reanalysis remains a frontier in climate research. ECMWF's coupled DA program, CERA-20C (Laloyaux et al. 2018), is now inactive, and ECMWF no longer hosts the output from CERA-20C. In this study, we will use "strongly" coupled DA, which (i) ensures that the coupled atmosphere-ocean-ice state is internally consistent and (ii) synthesizes observational and dynamical constraints across each component simultaneously.

To circumvent the computational obstacles associated with DA in fully coupled models, lightweight DA methods have been developed primarily for paleoclimate reconstruction. The "offline" DA method uses a static, uninformed prior from pre-existing model output (e.g., Hakim et al. 2016; Franke et al. 2017; Steiger et al. 2014, 2018; Samakinwa et al. 2021; Tierney et al. 2020; Osman et al. 2021; Valler et al. 2024; Smerdon et al. 2023). "Online" methods use a time-evolving prior that is informed by the previous initial conditions produced by data assimilation. Online DA requires integrating a forecast model after each assimilation step, and the expensive forecasting causes a computational bottleneck.

Data-driven approaches that emulate climate models can overcome the computational bottleneck. The linear inverse model (LIM) has been tested in annual-mean DA with proxies over the last millennium (Perkins and Hakim 2021) and for subseasonal forecasting (Hakim et al. 2022). LIMs have been applied to study dynamics and predictability of the El Niño–Southern Oscillation (ENSO) (e.g., Penland and Sardeshmukh 1995; Shin et al. 2021; Vimont et al. 2014; Lou et al. 2020; Kido et al. 2023), meridional modes (Vimont 2012), global surface temperatures (Newman 2013), SSTs in the North Atlantic (e.g., Zanna 2012) and North Pacific (e.g., Newman 2007; Newman et al. 2016; Zhao et al. 2024), the Pacific-North American pattern (Henderson et al. 2020), hydroclimate (Coats et al. 2020; Tseng et al. 2021), and sea ice (Brennan et al. 2023). LIMs are computationally efficient, enabling coupled assimilation of observations across Earth system components, e.g., pressure observations in the atmosphere and SST observations

in the ocean can each inform both SST and SLP in coupled DA. Combining LIMs with data assimilation presents an opportunity to constrain and quantify uncertainty in the historical climate record.

Here we use coupled data assimilation to reconstruct monthly and globally resolved SST, near-surface air temperature (T), SLP, and SIC over 1850–2023. The novelty of our approach compared to past reanalyses is that we constrain all climate fields simultaneously with (i) coupled dynamics and (ii) observations across climate components. Our DA method is made computationally tractable by efficient emulators (linear inverse models), which are trained on eight CMIP6 models and capture the essential dynamics for monthly reanalysis. We combine forecasts from LIMs with a Kalman filter to produce a coupled reconstruction with time-varying uncertainty quantification. Section 2 describes methods and data, including linear inverse models, data assimilation, validation with an out-of-sample pseudo-reconstruction, observations, and comparison datasets. Section 3 presents the historical reconstruction. Section 4 discusses the implications of the results for interpreting climate variability and changes and caveats of the method. Section 5 presents the conclusions.

2. Methods and data

In this section, we describe the reconstruction method, validation, and data sources. The reconstruction of monthly means consists of (i) a monthly forecast, for which we use LIMs that emulate eight CMIP6 models, and (ii) data assimilation in every month, for which we use the classic Kalman filter (Kalman 1960; Kalnay 2003). We validate the method with a pseudo-reconstruction of a climate model's 1850–2014 historical simulation (MPI-ESM1-2-HR), from which we draw observations that mimic the true observing network.

a. Linear inverse models

Anomalies around an equilibrium state in the nonlinear climate system can be approximated as a stochastically forced, linear dynamical system (e.g., Hasselmann 1976; Penland and Sardeshmukh 1995; Penland 1996):

$$\frac{d\mathbf{x}}{dt} = \mathbf{L}\mathbf{x} + \mathbf{S}\eta, \tag{1}$$

where $\mathbf{x}$ is a state vector of N principal components of SST, T, SLP, and SIC, $\mathbf{L}$ is an $N \times N$ linear operator representing the deterministic dynamics, and $\mathbf{S}\eta$ approximates the unresolvable nonlinear dynamics as stochastic forcing with an $N \times N$ noise-amplitude matrix, $\mathbf{S}$, and a vector, η, of independent, Gaussian white noise with unit variance and length N.

LIMs typically assume stationary statistics, but Shin et al. (2021) extend the LIM framework to include monthly variations in the dynamics. The monthly, or "cyclostationary" LIM, has been applied to ENSO (Shin et al. 2021; Vimont et al. 2022; Kido et al. 2023). We build on this recent work and use cyclostationary LIMs to model global SST, T, SLP, and SIC. We use the fixed-phase approach (OrtizBeviá 1997) to train the 12 $\mathbf{L}_j$ operators in the cyclostationary LIM, where j indicates the month:

$$\mathbf{L}_j = \tau^{-1} \log[\mathbf{C}_j(\tau)\mathbf{C}_j(0)^{-1}], \quad \text{for } j = 1, 2, ..., 12. \tag{2}$$

$\mathbf{C}_j(\tau)$ and $\mathbf{C}_j(0)$ are the τ-lag and zero-lag covariance matrices of $\mathbf{x}$ for month j, and $\tau = 1$ month in all of the following equations. The stochastic amplitude matrices, $\mathbf{S}_j$, are estimated from the fluctuation-dissipation relation of Equation (1) (Penland and Matrosova 1994),

$$\frac{d\mathbf{C}_j(0)}{dt} = \mathbf{L}_j\mathbf{C}_j(0) + \mathbf{C}_j(0)\mathbf{L}_j^T + \mathbf{Q}_j, \tag{3}$$

where $\mathbf{Q}_j = \mathbf{S}_j\mathbf{S}_j^T$. We follow Shin et al. (2021) in estimating the cyclostationary $\mathbf{Q}_j$ as

$$\mathbf{Q}_j = \frac{\mathbf{C}_{j+1}(0) + \mathbf{C}_{j-1}(0)}{2\Delta t} - [\mathbf{L}_j\mathbf{C}_j(0) + \mathbf{C}_j(0)\mathbf{L}_j^T], \tag{4}$$

with $\Delta t = 1$ month. Before computing $\mathbf{L}_j$ and $\mathbf{Q}_j$, we take the 3-month running means of $\mathbf{C}_j(\tau)$ and $\mathbf{C}_j(0)$, e.g., we estimate $\mathbf{C}_j(\tau) \approx \langle \mathbf{C}_{j-1}(\tau), \mathbf{C}_j(\tau), \mathbf{C}_{j+1}(\tau) \rangle$ (Shin et al. 2021). As in other LIM studies (e.g., Penland 1996), we remove any negative eigenvalues in $\mathbf{Q}_j$ and rescale remaining eigenvectors to conserve the original variance.

The LIM produces forecasts at lead $\tau = 1$ month from integrating (1) in time as

$$\mathbf{x}(t+\tau) = \mathbf{G}_j\mathbf{x}(t) + \mathbf{n}, \tag{5}$$

where $\mathbf{G}_j = \exp(\mathbf{L}_j\tau) = \mathbf{C}_j(\tau)\mathbf{C}_j(0)^{-1}$. The integrated stochastic term, $\mathbf{n}$, equals 0 in a deterministic forecast, such as the prior-mean forecast in the Kalman filter as described below.

The forecast equation for the error covariance, assuming no correlation between error and state, is

$$\mathbf{P}(t+\tau) = \mathbf{G}_j\mathbf{P}(t)\mathbf{G}_j^T + \mathbf{N}_j(\tau). \tag{6}$$

To solve for $\mathbf{N}_j(\tau)$, we extend the logic that applies to the stationary LIM (Hakim et al. 2022; Penland 1989) for the cyclostationary case. Equation (6) must be valid for any month's initial condition, including $\mathbf{C}_j(0)$, from which the monthly forecast must arrive at $\mathbf{C}_{j+1}(0)$ because the statistics are cyclostationary, therefore:

$$\mathbf{N}_j(\tau) = \mathbf{C}_{j+1}(0) - \mathbf{G}_j\mathbf{C}_j(0)\mathbf{G}_j^T. \tag{7}$$

We train separate LIMs to emulate the following eight CMIP6 models: CESM2, GFDL-ESM4, HadGEM3-

GC3.1-LL, SAM0-UNICON, UKESM1.0-LL, NorESM2-LM, EC-Earth3, and E3SM-2-0. For training, we use preindustrial-control simulations with the 1850–2014 historical simulations appended. Our selection of models is informed by Lou et al. (2023), which found that this subgroup performs best in an analog method for ENSO forecasting, although we make two changes: we remove HadGEM3-GC3.1-MM to prevent having two versions of HadGEM3, and we substitute E3SMv2.0 (Qin et al. 2024) for CIESM because of issues simulating sea ice in CIESM (Lin et al. 2020). LIM training is summarized in Appendix A. LIMs are trained separately for each model using monthly mean anomalies, and each LIM has a minimum of 665 years of training data (500 preindustrial and 165 historical years).

We regrid all training data to 2° resolution (96×144 latitude-longitude grid). For consistency with observations, which are expressed as anomalies relative to a 1961–1990 climatology, we remove the grand and climatological means calculated over 1961–1990 for each model. Separately for each model and state variable, we compute EOFs area-weighted by the square-root of the cosine of latitude for SST, T, SLP, Northern Hemisphere (NH) SIC, and Southern Hemisphere (SH) SIC. We retain approximately 85% of each field's variance in the truncated state. We form each model's standardized state vector from its principal components, $\mathbf{x}_k$, as:

$$\mathbf{x} = \begin{bmatrix} \mathbf{x}_{SST}/\sigma_{SST} \\ \mathbf{x}_{T}/\sigma_{T} \\ \mathbf{x}_{SLP}/\sigma_{SLP} \\ \mathbf{x}_{SIC_{NH}}/\sigma_{SIC_{NH}} \\ \mathbf{x}_{SIC_{SH}}/\sigma_{SIC_{SH}} \end{bmatrix},$$

where σ_k^2 is the retained variance after EOF truncation of field k. We use the standardized state vectors $\mathbf{x}$ to compute covariance matrices for each model, and we project into and out of the LIM basis by storing the EOFs and scale factors, σ_k, for each field. Each LIM is run independently in parallel through the data assimilation framework.

b. Data assimilation

Given a prior forecast of the state's monthly mean $\mathbf{x}_f$ and error covariance $\mathbf{P}_f$, we assimilate observations to produce the posterior analysis $\mathbf{x}_a$ and $\mathbf{P}_a$ using the Kalman filter:

$$\mathbf{x}_a = \mathbf{x}_f + \mathbf{K}(\mathbf{y} - \mathbf{H}\mathbf{x}_f), \tag{8}$$

$$\mathbf{P}_a = [\mathbf{I} - \mathbf{K}\mathbf{H}]\mathbf{P}_f, \tag{9}$$

$$\mathbf{K} = \mathbf{P}_f\mathbf{H}^T[\mathbf{H}\mathbf{P}_f\mathbf{H}^T + \mathbf{R}]^{-1}, \tag{10}$$

where $\mathbf{K}$ is the Kalman gain, $\mathbf{y}$ is the vector of observations, $\mathbf{H}$ is the linear observation operator, and $\mathbf{R}$ is the observation error covariance. After solving (8–10) for a given month, we forecast the next month from (5), with $\mathbf{n} = 0$, and (6).

Our method is "strongly coupled online DA," where "strongly coupled" means that we assimilate observations concurrently across the atmosphere–ocean–ice system, and all fields influence each other through cross-component covariances. "Online" means that we use a forecast model with the previous assimilation step's initial conditions to inform the prior. Because this method uses the classic Kalman filter and propagates $\mathbf{P}_f$ exactly, we avoid the sample error and localization issues that arise when estimating $\mathbf{P}_f$ in an ensemble Kalman filter (Evensen 1994; Houtekamer and Zhang 2016). However, ensemble-member trajectories are needed to analyze statistics of temporal variability, and this variability must be constrained by dynamics rather than sampled independently (Emile-Geay et al. 2024). We solve this problem with a modified version of the ensemble Kalman filter, described subsequently, that has no impact on the mean or covariance, (8–9), but rather simply provides sample estimates from the posterior distribution.

We generate ensemble members using the perturbed-observations version of the ensemble Kalman filter (Houtekamer and Mitchell 1998; Burgers et al. 1998), except we use the exact prior covariance forecast from the classic Kalman filter (6, 9). For each LIM, we initialize 200 ensemble members in January 1850 with random draws from a multivariate-normal distribution with covariance $\mathbf{C}_1(0)$. Each ensemble member is updated using (8), with $\mathbf{x}_f^n$ corresponding to ensemble member n in place of the ensemble mean, and $\mathbf{y}^n$ is a multivariate-normal random draw of the observations with mean $\mathbf{y}$ and covariance $\mathbf{R}$. After the assimilation, each $\mathbf{x}_a^n$ is advanced to the next month using (5). The noise term in (5), $\mathbf{n}$, becomes a random draw from $\mathbf{N}_j(\tau)$ in (7) for each ensemble member. Because our LIMs are built to forecast monthly means, we can draw from the distributions of the monthly statistics rather than stochastically integrating (Penland and Matrosova 1994) each ensemble member.

An additional benefit of the ensemble is that we can propagate temporally correlated observation errors that are associated with uncertainties in bias corrections. For example, HadSST4 (described below) provides a 200-member ensemble of monthly SST observations to represent temporally correlated errors (Kennedy et al. 2019). To incorporate these errors, we let $\mathbf{y}$ vary across the ensemble members, but each of our 200 ensemble members $\mathbf{x}^n$ is paired at every timestep with the corresponding ensemble member n from the HadSST4 ensemble.

c. Observations

We use four sources of observations corresponding to each of the four state variables (SST, T, SLP, SIC). All

observations are anomalies relative to a 1961–1990 climatology, which is the period chosen by Kennedy et al. (2019) and Osborn et al. (2021).

SST observations are from HadSST4 version 4.0.1.0 (Kennedy et al. 2019), provided by the Met Office Hadley Centre on a $5° \times 5°$ grid. HadSST4 quality controls and corrects biases in the in situ measurements from ICOADS 3.0.0 (1850–2014) and ICOADS 3.0.1 (2015–present), the central database of ship records (Freeman et al. 2017). HadSST4 provides non-infilled data as monthly means spanning 1850–present, and ship coverage varies substantially over time (Figure 1). Measurement and sampling errors are provided for every gridcell and month with data, and error covariance matrices are provided that estimate the spatially correlated errors. We include these sources of error in $\mathbf{R}$. Temporally correlated errors from uncertain bias corrections are estimated with a 200-member ensemble of observations, and we account for these errors with our ensemble DA method, described in Section 2b.

Observations of near-surface air temperature (T) over land are from CRUTEM5 version 5.0.2.0 (Osborn et al. 2021). Weather-station data is quality controlled, bias-corrected, and provided as monthly means with error estimates on a $5° \times 5°$ grid. We include CRUTEM5's time-varying measurement and sampling errors in $\mathbf{R}$.

Marine SLP observations are from ICOADS Enhanced Release 3.1 for 1850–2014 and Release 3.0.2, for 2015–2023 (Freeman et al. 2017). SLP data are provided as monthly means on a $2° \times 2°$ grid, along with the number of observations, n_{obs}, in each month and the intra-month standard deviation, s, of the observations in each gridcell. The baseline climatology for anomalies is from Hersbach et al. (2020). There are a large number of SLP observations due to the finer grid of ICOADS compared to HadSST4. We eliminate observations with $n_{\text{obs}} < 5$, which are expected to have a low signal-to-noise ratio. For months that have data in more than 3000 gridcells, we mask up to 40% of the values between 25°S and 60°N using random sampling. These limits increase computational efficiency of the assimilation and maintain a reasonable balance between the number of SLP and SST observations, otherwise there would be approximately five times as many SLP as SST observations. Valler et al. (2024) also reduce the number of ICOADS observations of SLP in their atmospheric reanalysis and set a similar threshold of $n_{\text{obs}} = 10$ per gridcell. Past studies identified a bias in ICOADS SLP data before 1870, which is discussed in Slivinski et al. (2019), Freeman et al. (2017), and Allan and Ansell (2006). NOAA 20CRv3 performed a bias correction of the pre-1870 SLP observations, so we substitute the 1850–1870 SLP from ICOADS with the collocated values from NOAA 20CRv3. ICOADS does not provide an estimate of measurement and sampling errors which comprise the diagonal terms in $\mathbf{R}$. As described in Appendix B, we estimate $\mathbf{R}$ from the intra-month spread in individual observations and the variance across neighboring observations.

Sea ice observations are provided by the NOAA/NSIDC Climate Data Record (CDR) of Passive Microwave Sea Ice Concentration, Version 4 from 11/1978–09/2023 and Near-Real-Time, Version 2, for 10/2023–12/2023 (Meier et al. 2021b,a). We coarsen the observations from 25 km to 2° resolution. At each timestep in the assimilation with satellite data, we use a subset of the available data, which has nearly complete coverage of the polar regions. We retain all observations with SIC ranging from 0.01 to 0.98, and we retain 40% of the remaining observations using random sampling. For measurement and sampling errors that form the diagonal terms in $\mathbf{R}$, we use the provided standard deviations of daily values, but we set the minimum error to 1 percentage point. As described in Appendix B for SLP, these intramonth standard deviations approximate the monthly mean error. For SIC, they are calculated across both the NASA Team and Bootstrap algorithms, sampling the systematic error across data-processing methods. Errors are small in open water and pack ice but are often between 30 and 50 percentage points in partial ice cover. We do not have satellite data for sea ice from 1961–1978, but we require a full climatology from 1961–1990 to calculate the SIC anomalies relative to a baseline that is consistent with the HadSST4 anomalies. The mean of the eight models used for LIM training agrees well with observations over the satellite era (SI of Roach et al. 2020; Notz and SIMIP Community 2020), so we combine the multi-model mean of the eight historical simulations from 1/1961–11/1978 with the satellite data from 12/1978–12/1990, and we use the merged climatology from 1961–1990 as the reference for SIC anomalies. Because solutions to (8) are not restricted to SIC between zero and one, we use the climatology in postprocessing to ensure that SIC is between zero and one.

d. Validation: Pseudo-reconstruction of an out-of-sample model

To test our method, we mimic the real reconstruction problem and attempt to reconstruct the 1850–2014 historical simulation from a climate model. Our target model is MPI-ESM1-2-HR, ensemble member r1i1p1f1 (Mauritsen et al. 2019). We have chosen MPI-ESM1-2-HR because it is a difficult test of the method given that, unlike nearly all other models, it has cooling in the Southern Ocean from 1980–2014. It also has a low-bias in Antarctic sea ice (Roach et al. 2020) and substantially different ENSO statistics and radiative feedbacks (Bloch-Johnson et al. 2024) compared to the models used for LIMs and priors in the data assimilation. The pseudo-reconstruction's target is out-of-sample because MPI-ESM1-2-HR is not used for LIM training; the dynamics of the target model are unknown to our eight forecast models.

We draw pseudo-observations from the target simulation at the same times and locations where real observations are available for SST, T, SLP, and SIC. Random errors are added to the pseudo-observations by sampling from the real observation errors in **R**. Note that real observations also have biases and unknown, unquantified errors which make the real reconstruction more challenging than this test. On the other hand, the LIMs used as model priors are selected based on their ability to collectively emulate reality rather than the target model of the pseudo-reconstruction.

Figure 2 shows timeseries representing climate variability from the pseudo-reconstruction. The ensemble mean is calculated as the grand mean across all 1600 ensemble members (8 LIMs × 200 members), and the ensemble shading spans the 17th–83rd percentiles. We note that any one of the eight LIM-DA systems may not have a posterior distribution that spans the true state by itself. However, the grand ensemble of posterior distributions from all eight LIM-DA systems, which includes the spread from model error, generally spans the truth (Figure 2). Because the grand ensemble represents eight separate DA systems, its distribution is non-Gaussian.

The metrics in Figure 2 are calculated as follows, with anomalies representing the departures from the 1961–1990 climatological annual cycle unless stated otherwise:

- Atlantic multidecadal variability (AMV) is the monthly mean SST anomaly in the North Atlantic (0°–60°N, 80°W–0°W) minus the global mean; the mean of the index from 1900–1970 is removed before plotting (Trenberth and Shea 2006).

- The Pacific Decadal Oscillation (PDO) is the leading EOF of the monthly mean SST anomaly in the North Pacific (20°–70°N) after removing the global mean (Newman et al. 2016).

- Nino3.4 is the monthly mean SST anomaly over 170°W to 120°W and 5°S–5°N, with the 30-yr running mean removed.

- The zonal SST gradient in the tropical Pacific is the mean SST anomaly in the west (80°E–150°E) minus the east (160°W–80°W), spanning 5°S–5°N (e.g., Heede and Fedorov 2023).

- SST$^{\#}$, which denotes the tropical SST contrast, is the mean of the warmest 30% of all tropical SSTs (30°S–30°N) minus the mean tropical SST, and the 1961–1990 mean is removed (Fueglistaler 2019). We note that SST$^{\#}$ requires actual SSTs, not just anomalies. To estimate SST$^{\#}$ for the pseudo-reconstruction, we add reconstructed anomalies to the target model's 1961–1990 climatology, which assumes outside knowledge of the target's climatology. For the actual reconstruction of SST$^{\#}$, we add the reconstructed anomalies to the 1961–1990 climatology from HadISST2.1's ensemble mean. We show the 5-year running mean of SST$^{\#}$ for consistency with Fueglistaler and Silvers (2021).

- Southern Ocean SST is the zonal-mean SST anomaly from 50°–70°S (Doddridge and Marshall 2017).

- Global-mean near-surface air temperature (GMSAT) is the global-mean T anomaly.

- The Walker circulation, measured by the zonal SLP gradient, is the mean SLP anomaly in the west Pacific (130°E–150°E) minus the central-east Pacific (160°W–120°W), spanning 5°S–5°N (e.g., Heede and Fedorov 2023).

- The Southern Annular Mode (SAM) is the standardized zonal-mean SLP anomaly at 40°S ± 2° minus the standardized zonal-mean SLP anomaly at 65°S ± 2° (Gong and Wang 1999); the reference period for standardization is 1961–1990, and each month is standardized separately.

- Sea ice area is the sum of the products of SIC and gridcell area, and we note that a common land mask is used when comparing ice area across datasets.

Most large-scale metrics are reconstructed with accuracy. We assess performance by the Pearson correlation (R), the fraction of variance explained (R^2), and the Nash-Sutcliffe Efficiency (NSE),

$$\mathrm{NSE} = 1 - \frac{\sum(x_i - \hat{x}_i)^2}{\sum(x_i - \bar{x})^2},$$

which accounts for the relative phasing of the true timeseries (x_i) versus the reconstructed timeseries ($\hat{x_i}$), the signal amplitude, and bias. The NSE has an upper bound equal to one and can become negative from biases in the mean or amplitude of variability (Nash and Sutcliffe 1970). We find $R^2 > 0.80$ for the AMV, PDO, Nino3.4, the 30-year rolling 1σ of Nino3.4, the zonal SST gradient in the tropical Pacific, GMSAT, Southern Ocean SST, the Walker circulation (zonal SLP gradient), the SAM, and Arctic ice area. The tropical SST contrast, $SST^{\#}$, has the lowest R^2 at 0.31. NSE values are also shown in Figure 2.

The reconstruction of the Walker circulation has a damped amplitude compared to the target, which is due to the EOF truncation of SLP in the LIM training. We show an additional version of the target model's Walker circulation, which is calculated after truncating the target's SLP into the leading 30 EOFs. Truncation is expected to affect tropical SLP because the variance in tropical SLP is low compared to the variance at higher latitudes, but truncation does not appear to have a substantial influence on other metrics.

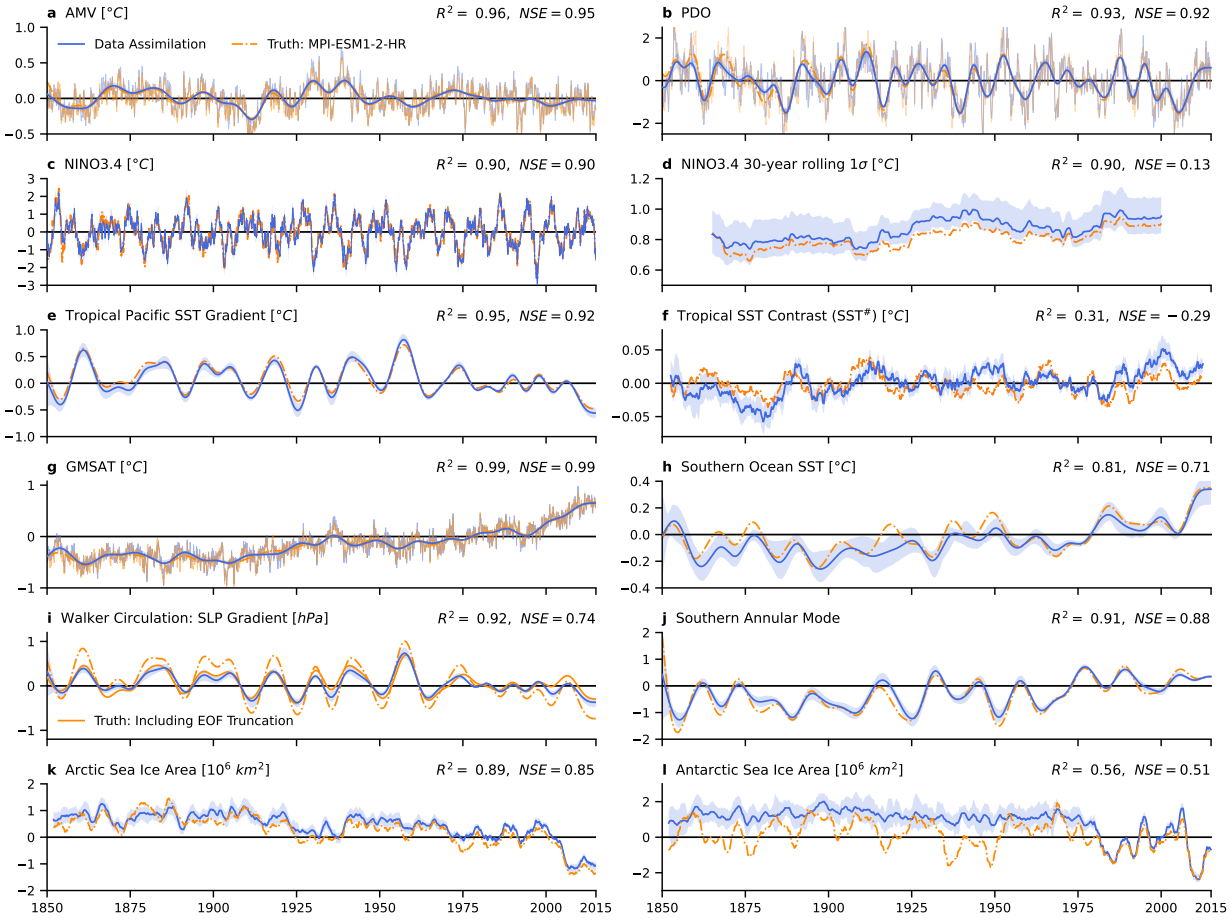

FIG. 2. **Validation by pseudo-reconstruction: timeseries.** (Orange) True values from the target model, the 1850–2014 historical simulation from MPI-ESM1-2-HR. (Blue) Result from data assimilation, showing mean of 1600 ensemble members; shading denotes ensemble 17$^{\text{th}}$ and 83$^{\text{rd}}$ percentiles, i.e., *likely* range. **(a)** Atlantic Multidecadal Variability with 10-yr low-pass filter and monthly values as thin lines. **(b)** Pacific Decadal Oscillation with 6-yr low-pass filter and monthly values as thin lines. **(c)** Monthly Nino3.4 with 30-yr running mean removed. **(d)** Rolling 30-yr standard deviation of Nino3.4 in panel **c**. **(e)** Zonal gradient of tropical Pacific SST with 10-yr low-pass filter. **(f)** Tropical SST contrast, SST$^{\#}$, 5-yr running mean. **(g)** Global-mean near-surface air temperature (GMSAT) with 10-yr low-pass filter and monthly values in thin lines. **(h)** Zonal mean of Southern Ocean SST (50°–70°S) with 10-yr low-pass filter. **(i)** Walker circulation, i.e., zonal SLP gradient across tropical Pacific, with 10-yr low-pass filter. **(j)** Southern Annular Mode with 10-yr low-pass filter. **(k)** Total area of Arctic sea ice with 24-month running mean. **(l)** Total area of Antarctic sea ice, with 24-month running mean. Calculation of metrics is described in Methods Section 2d.

Antarctic sea ice has $R^2 = 0.56$ and is biased high in the reconstruction before 1979. The reason for this bias is that the target model is biased low relative to the multi-model mean of the LIMs and relative to the satellite record (Roach et al. 2020). There are decadal periods of abrupt ice loss in the target model which are not captured in the reconstruction. These ice-loss events are associated with brief warming episodes in Southern Ocean SST (Figure 2h), which are also not detected in the reconstruction. While we do not know whether such ice-loss events happen in nature, it is worth noting that if they do occur, our method cannot identify them in the sparse instrumental observations. Despite missing these decadal warmings, the lower-frequency variability in Southern Ocean SST and the SAM is captured by the reconstruction.

Figure 3 shows the pattern of trends in annual-mean SST for 1900–1979 and 1980–2014. Local trends are divided by the global-mean trend to emphasize the patterns. We also show the reconstruction's ensemble spread (1σ) in trend patterns, which highlights regions of elevated uncertainty. It is important to recall that observations in the Southern Ocean and southeast Pacific are sparse even after 1980 (Figure 1c), which is evident in our uncertainty quantification.

To further illustrate the uncertainty, we show trends from individual ensemble members (Figure 3c,g). These ensemble members show more cooling in the Southern Ocean