Data assimilation for networks of coupled oscillators: Inferring unknown model parameters from partial observations

Lauren D. Smith[1,*] and Georg A. Gottwald[2,†]

[1] *Department of Mathematics, The University of Auckland, Auckland 1142, New Zealand*
[2] *School of Mathematics and Statistics, The University of Sydney, Sydney, NSW 2006, Australia*

(Dated: September 7, 2023)

We develop a data-driven method combining a forecast model with unknown parameters and partial noisy observations for networks of coupled oscillators. Employing network-specific localization of the forecast covariance, an Ensemble Kalman Filter with state space augmentation is shown to yield highly accurate estimates of both the oscillator phases and unknown model parameters in the case where only a subset of oscillator phases are observed. In contrast, standard data assimilation methods yield poor results. We demonstrate the effectiveness of our approach for Kuramoto oscillators and for networks of theta neurons.

Many natural phenomena and engineering applications can be described as networks of coupled oscillators, for example, the firing of neurons in the brain[1–3] and the dynamics of power grids[4–7]. However, mathematical models for these systems are often incomplete, with unknown parameters and simplifying physical assumptions. Scientists are now seeking data-driven methods to estimate unknown parameters and provide an improved estimate of the dynamic state from partial observations.

A pertinent question is: Can one accurately estimate all the oscillators' phases as well as unknown model parameters, such as oscillator natural frequencies, if only a subset of the oscillator phases are observed? For example, in a power grid, we may be able to observe all the power stations, but none of the consumers. The framework of data assimilation (DA) provides a unifying framework that has the potential to achieve this by estimating a system's state and unknown parameters through a combination of uncertain forecasts from a given model and noisy partial observations.

We focus specifically on the Ensemble Kalman Filter (EnKF)[8–14] as a DA method. Although the framework of Kalman filters can strictly only be applied to linear systems, the EnKF has been successfully used in numerous non-linear settings[15]. In particular, the Kalman filter has been widely used in the weather forecasting community combining model forecasts and noisy observational data to yield improved forecasts[8–14]. To date, there are no applications of DA to the broad class of non-linear systems that model coupled oscillators on complex networks. Indeed, we find that the standard EnKF algorithm performs poorly for networks of coupled oscillators unless very large ensembles are used. However, we show that when a novel localization method is incorporated, specific to dynamics on networks, the method is highly effective at estimating both oscillator phases and unknown model parameters, even when a small ensemble is used and only a fraction of the phases are observed. This sets DA apart from other data-driven parameter estimation methods that require all oscillator phases to be observed[16–19]. We demonstrate our method for two models of coupled oscillators; 1. The Kuramoto model[20,21] which is a simplification of power grid dynamics and captures the synchronization phenomena observed in many coupled oscillator systems, and 2. Networks of theta neurons[22–25] which model neural dynamics of the brain.

Kalman filters aim to estimate the posterior distribution of the state given a model forecast and noisy observations. The distributions are parametrized by the mean and the covariance of the variables, assuming a multivariate Gaussian distribution. In ensemble Kalman filters the mean and covariances are approximated by an ensemble of independent forecasts. In particular, at each discrete time step t_k, the mean (called the "analysis") is provided by

$$\bar{X}^{\mathrm{a}}_k = \bar{X}^{\mathrm{f}}_k + K_k(Y_k - H\bar{X}^{\mathrm{f}}_k),$$

where Y_k is the vector of observations, H is the observation matrix that projects from state space to observation space, $\bar{X}^{\mathrm{f}}_k$ is the ensemble mean of the model forecast, obtained by simulating the ensemble with the full nonlinear model from time t_{k-1} to t_k initialized with the previous analysis ensemble with mean $\bar{X}^{\mathrm{a}}_{k-1}$ and covariance $P^{\mathrm{a}}_{k-1} = (I - K_k H)P^{\mathrm{f}}_{k-1}$, and K_k is the Kalman gain matrix

$$K_k = P^{\mathrm{f}}_k H^T (H P^{\mathrm{f}}_k H^T + R)^{-1}, \tag{1}$$

where R is the observational error covariance matrix. The forecast error covariance matrix P^{f}_k is estimated at each time step from the ensemble as a Monte-Carlo approximation. It is this estimation of the error covariance matrix via the ensemble which is propagated using the full nonlinear dynamics which sets ensemble methods apart from classical Kalman filters which often assume a constant background error covariance matrix. In order to estimate unknown model parameters in addition to the state space variables, we employ state space augmentation[11,12,26]. Letting $X = (\phi_1, \ldots, \phi_n)$ denote the state variables and $\xi_1, \ldots, \xi_m$ denote the unknown parameters, the new state variables are $\tilde{X} =$

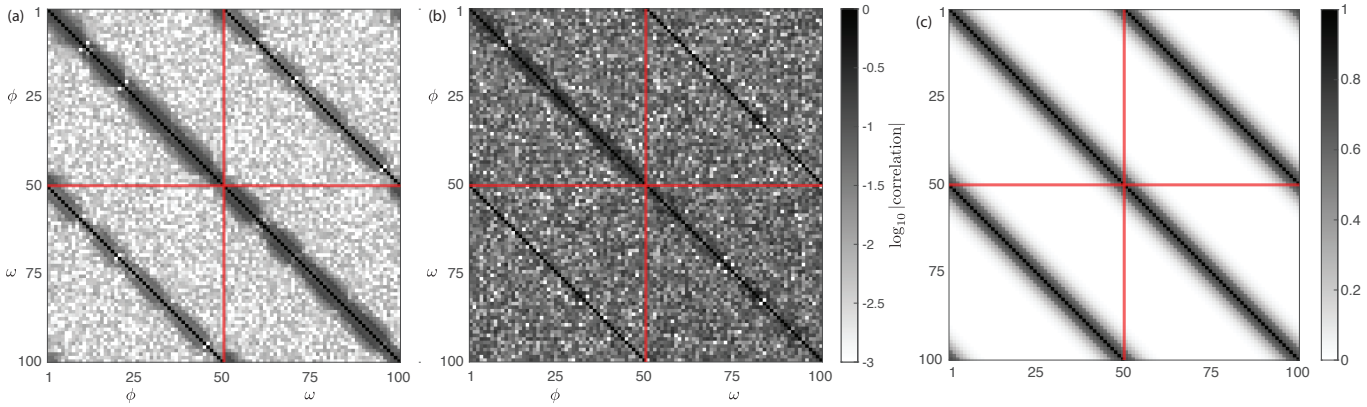

FIG. 1. (a) Average correlation matrix Q for a large ensemble ($M = 10100$), and (b) for a small ensemble ($M = 101$), obtained by applying the standard EnKF to the Kuramoto model (4). The augmented state space is $(\phi_{1,\ldots,50}, \omega_{1,\ldots,50})$, hence the block matrix substructure (indicated by red lines) in Q with ϕ-ϕ correlations, ϕ-ω correlations, ω-ϕ correlations and ω-ω correlations. The Kuramoto model parameters are $\kappa = 27$ and $\omega_i \sim \mathcal{N}(0, 0.1)$. (c) The localization matrix $\mathcal{L}$ (6) for the ring topology with $N = 50$ and $r = 3$ using $\lambda = 0.46$.

$(\phi_1, \ldots, \phi_n, \xi_1, \ldots, \xi_m)$, with dynamics $\dot{\xi}_j = 0$ (see Appendix A for full details).

For networks of coupled oscillators, the state variables are typically phases which are 2π-periodic. As such, care needs to be taken when computing ensemble means and covariances. Consider an ensemble of M N-dimensional phase vectors, $\boldsymbol{\phi}^1, \ldots, \boldsymbol{\phi}^M$, with $\boldsymbol{\phi}^k = (\phi_1^k, \ldots, \phi_N^k)$. We assume that for each $j = 1, \ldots, N$, the ensemble ϕ_j^k, $k = 1, \ldots, M$, is relatively concentrated to one region of the unit circle. The ensemble mean $\bar{\phi}_j$ for each $j = 1, \ldots, N$ (and, hence $\bar{X}^{\mathrm{f}}$) is computed as the mean in the complex plane, i.e.,

$$\rho e^{i\bar{\phi}_j} = \frac{1}{M} \sum_{k=1}^{M} e^{i\phi_j^k}. \tag{2}$$

The ensemble forecast error covariance matrix is computed as

$$P_{jl}^{\mathrm{f}} = \frac{1}{M-1} \sum_{k=1}^{M} F\left(\phi_j^k - \bar{\phi}_j\right) F\left(\phi_l^k - \bar{\phi}_l\right), \tag{3}$$

for $j, l = 1, \ldots, N$, where $F(\theta) = \mathrm{mod}(\theta + \pi, 2\pi) - \pi$ is used to ensure that the differences are centered at 0 and are in the interval $[-\pi, \pi)$.

Since the distributions of X^{a} and X^{f} are estimated by ensembles, there is a trade-off between the ensemble size and computational speed. Large ensembles are desired to accurately estimate the covariances P_k^{f}, but small ensembles enable faster computation. For small ensembles, spurious correlations can have a significant effect at each DA step. To suppress the effect of spurious correlations so-called localization is employed, whereby geographical/topological information is used to predetermine which correlations will be spurious[27]. In standard applications, such as weather forecasting, correlations are deemed spurious if the grid points are far apart. It is less clear how to best incorporate localization for dynamics on networks. The following examples illustrate the need for localization in networks of coupled oscillators, as well as how localization should be performed.

Example 1: The Kuramoto model. To demonstrate DA, spurious correlations, and appropriate localization for networks of coupled oscillators, we begin with the Kuramoto model[20,21]. For N oscillators, the dynamics of the i-th oscillator is

$$\dot{\phi}_i = \omega_i + \frac{\kappa}{N} \sum_{j=1}^{N} A_{ij} \sin(\phi_j - \phi_i), \tag{4}$$

where ω_i is the oscillator's natural frequency, κ is the coupling strength and A is the network adjacency matrix. This model approximates the dynamics of many engineering, biological and chemical processes with interacting oscillators[21,28], for example, it models the power grid wherein the nodes represent the electrical phases of power stations and consumers, and the network edges represent power lines[5–7].

We return to the question posed in the introduction: Can DA accurately estimate all the phases ϕ_i and all the natural frequencies ω_i if only a subset of the phases are observed? We assume that κ and A are known.

Before addressing the above problem, we begin with a simpler problem that illustrates spurious correlations and how they can be mitigated: Can DA accurately estimate all the phases ϕ_i and all the natural frequencies ω_i if *all* of the phases are observed? For illustrative purposes we assume a ring network topology, such that nodes are arranged in a circle, and each node is connected to its $2r$ nearest neighbors (coupling radius r). To demonstrate spurious correlations, we record the covariance matrices P_k^{f} at each step of the DA process, and normalize them

to generate correlation matrices Q_k. For the DA process, we generate noisy observations Y_k by first simulating the model and then adding noise with prescribed covariance $R = \eta^2 I$ (here $\eta = 0.02$). Averaging the Q_k over the whole DA process yields the average correlation matrix Q, shown in Fig. 1(a,b) for both a large ensemble size $M = 100(2N + 1) = 10100$ and a small ensemble size $M = 2N + 1 = 101$. Since the augmented state space is $2N$-dimensional, the minimum ensemble size to ensure a full rank covariance matrix is $M = 2N+1$. For the large ensemble (Fig. 1(a)), we observe a rapid decay in correlation away from the leading diagonal in each sub-block matrix, which reflects the network connectivity of the underlying ring topology. On the other hand, for the small ensemble (Fig. 1(b)), there is a high correlation region close to the diagonal, but there are also large spurious correlations far from the diagonal, i.e., large correlations between nodes that are poorly connected. Moreover, the information about the network topology appears to be completely lost in the ϕ-ω correlations, which are crucial for estimating the ω_i (see Appendix A). To suppress these spurious correlations for small ensembles, we propose using localization: At each DA step, the covariance matrix P_k^{f} is replaced by

$$\tilde{P}_k^{\mathrm{f}} = \mathcal{L} \circ P_k^{\mathrm{f}}$$

where $\circ$ denotes the Schur (elementwise) product, and the localization matrix $\mathcal{L}$ has the following properties:

- $\mathcal{L}$ encodes the network connectivity, i.e., $\mathcal{L}_{ij}$ is large if nodes i and j are well connected, and small otherwise,
- $\mathcal{L}$ is a correlation matrix, i.e., $\mathcal{L}_{ii} = 1$ and $\mathcal{L}$ is positive semi-definite.

Consider the matrix exponential $\mathcal{E} = \exp(\lambda A)$, where λ is a tunable parameter and A is the network adjacency matrix (assumed to be symmetric). Since

$$\mathcal{E} = \exp(\lambda A) = \sum_{n=0}^{\infty} \frac{\lambda^n}{n!} A^n$$

is a weighted sum, with decreasing weights, of powers A^n, each counting the number of paths of length n, $\mathcal{E}$ encodes the network topology. Moreover, $\mathcal{E}$ is guaranteed to be a symmetric positive semi-definite matrix, i.e., a covariance matrix. $\mathcal{E}$ can be converted to a correlation matrix by renormalizing, i.e.,

$$L = \left(\mathrm{diag}(\mathcal{E})\right)^{-\frac{1}{2}} \mathcal{E} \left(\mathrm{diag}(\mathcal{E})\right)^{-\frac{1}{2}}, \tag{5}$$

where $\mathrm{diag}(\mathcal{E})$ denotes the matrix of diagonal elements of $\mathcal{E}$. Thus, L is a correlation matrix that encodes the network topology, fulfilling our requirements of a localization matrix. L would be sufficient if we only aimed to estimate phases via DA, for the augmented state space with both phases and frequencies, we use the block matrix

$$\mathcal{L} = \begin{pmatrix} L & L \\ L & L \end{pmatrix}. \tag{6}$$

An example of $\mathcal{L}$ is shown in Fig. 1(c) for the ring topology with $N = 50$ and $r = 3$ using $\lambda = 0.46$. The parameter λ controls the rate of decay for distant nodes and is chosen to suppress spurious correlations (see Appendix B for full details on the choice of λ). We will show that applying this localization significantly reduces the impact of spurious correlations on the DA process. We note that the matrix exponential for the ring topology has a similar profile to standard localization functions[27,29], however, the standard localization functions do not yield correlation matrices when applied to networks, as they rely on Euclidean geometry (see Appendix D for full details).

Returning to the more difficult problem of determining the ϕ_i and ω_i from limited observations, we consider the case where only 70% (35 out of 50) of the phases are observed. We find, remarkably, that even with this limited information, DA with localization does an excellent job at determining all of the phases and all of the unknown parameters ω_i, whereas the standard EnKF algorithm without localization performs poorly. This is shown in Fig. 2. Fig. 2(a) shows the root mean square (RMS) errors between the true values and the DA estimated values for both the phases ϕ_i (dashed) and the frequencies ω_i (solid), using both the standard EnKF algorithm (blue) and the EnKF algorithm with localization using (6) (red). The RMS error for the phases at time t is

$$E_\phi(t) = \frac{1}{\sqrt{N}} \sqrt{\sum_{i=1}^{N} \left(\phi_i(t) - \hat{\phi}_i(t)\right)^2}, \tag{7}$$

where $\hat{\phi}_i(t)$ is the DA estimated value and $\phi_i(t)$ is the true value. The RMS error for the frequencies, E_ω, is computed analogously. We see that localization greatly reduces the error in both the phases and the frequencies compared to the standard algorithm, with almost 10 times less error in each. The normalized residuals (differences between the true and estimated values, normalized by the maximum value) at the final time $t = 15$ for the respective DA approaches are shown in Fig. 2(b,c) for ϕ_i and ω_i, respectively. As well as observing that localization (red circles) greatly reduces the error compared to the standard algorithm (blue triangles), we see that for each approach (standard or localized) DA performs, as expected, better for the nodes that are observed (filled circles/triangles) compared to the unobserved nodes (open circles/triangles). Comparing the median residuals (solid and dashed horizontal lines in Fig. 2(b,c)), localization without node observation (dashed red line) yields a smaller median residual than the standard algorithm with node observation (solid

blue line), which again highlights the importance of localization. We note that DA struggles the most in the region with a large number of unobserved nodes, i.e., in the range $i = 20$ to $i = 40$. This is expected because there is less information available to unobserved nodes, and this is compounded if unobserved nodes are clustered together and their estimates are not sufficiently controlled by incoming observations.

We have used the ring topology here to illustrate spurious correlations and network-specific localization. Our methodology readily extends to general networks, and network-specific localization yields a great improvement in estimation accuracy (see Appendix C for a discussion of the localization method applied to Erdős-Renyi graphs).

In Appendix E we discuss the effect of changing the fraction of observed nodes. As expected, the estimation accuracy increases with the fraction of observed nodes. Localization amplifies this effect, and allows for a much smaller set of observed nodes to achieve the same estimation accuracy compared to the standard approach.

We remark that synchronization presents a challenge for data-driven problems in general when applied to models similar to the Kuramoto model. For sufficiently large coupling strengths κ, the phases will tend to a synchronous rotating state, which is effectively a stationary state in a rotating reference frame. Once synchronization occurs, observational data becomes degenerate, as no new information is being fed to the DA process. In the examples here, we show results for the transient period before synchronization occurs.

We now present a second example in which synchronization does not generally occur, and which is used to model activity within the brain:

Example 2: Networks of theta neurons[22–25]. For N neurons, the dynamics of the i-th neuron is given by

$$\dot{\phi}_i = 1 - \cos\phi_i + (1 + \cos\phi_i)\,(\zeta_i + \kappa I_i)\,, \qquad (8)$$

where κ is a coupling strength, ζ_i is an intrinsic parameter that determines the neuron's propensity to fire, and I_i is the input from other neurons to neuron i given by

$$I_i = \frac{2\pi}{N}\sum_{j=1}^{N} B_{ij} P(\phi_j),$$

where $P(\phi) = a(1-\cos\phi)^2$, with a such that $\int_0^{2\pi} P(\phi)d\phi = 2\pi$, and B_{ij} is the connectivity between nodes i and j, which may be negative for inhibitive coupling. When a neuron receives sufficient input from its neighbors, it fires, completing a revolution of the unit circle. This firing provides input to the neuron's neighbors which in turn may cause them to fire. In such a network, many irregular and sustained firing patterns are possible depending on the model parameters and the initial condition.

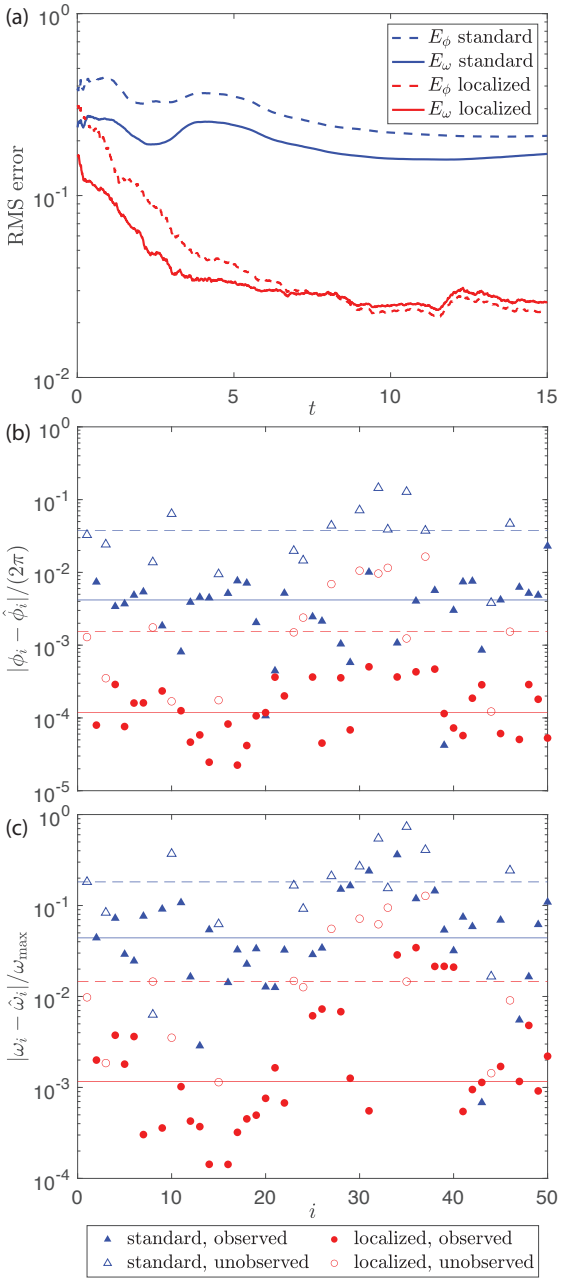

FIG. 2. (a) RMS errors (7) in ϕ (dashed) and ω (solid) for DA applied to the Kuramoto model (4) with 35 out of 50 phases observed. Results are shown for the standard EnKF with state space augmentation (blue) and EnKF with state space augmentation and network-based localization (6) with $\lambda = 0.46$ (cf. Fig. 1(c)) (red). (b), (c) Normalized residuals of ϕ_i and ω_i, respectively, at final time $t = 15$ for the standard EnKF algorithm (blue triangles) and for the EnKF algorithm with localization (red circles). Filled symbols indicate observed nodes, and open symbols indicate unobserved nodes. Medians of the residuals are shown as horizontal lines; blue for standard, red for localized, solid for observed, dashed for unobserved. Note the logarithmic scale on the vertical axes. The Kuramoto model parameters are $\kappa = 10$ and $\omega_i \sim \mathcal{N}(0, 0.1)$ with the ring topology with $r = 3$.

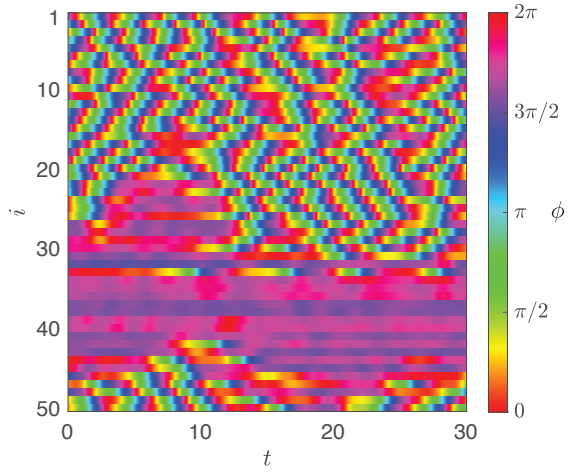

FIG. 3. Time evolution of the phase variables ϕ_i of the theta neuron model (8) showing a bump state. The theta neuron model parameters are $N = 50$, $\kappa = 2$ and $\zeta_i \sim \mathcal{N}(-0.4, 0.1)$, with connectivity matrix B described in the main text.

We consider a similar problem to that posed for the Kuramoto model: Can DA accurately estimate all the phases ϕ_i and all the intrinsic firing parameters ζ_i if only a subset of the phases are observed?

For the connectivity matrix B, we consider a similar topology to the ring topology, such that nodes are arranged in a ring and positively coupled with $B_{ij} = 1$ if they are within a coupling radius $r = 3$. We also include negative (inhibitory) long range coupling, such that $B_{ij} = -0.4$ for the three furthest nodes from each node. All other entries of B are zero. Similar connectivity assumptions are often made in modeling studies[30], and typically produce "bump states" such as that shown in Fig. 3, with a region of neurons that fire approximately periodically and another region of neurons that do not fire at all. Such bump states are thought to be connected to short term memory[31].

As for the Kuramoto model, correlation matrices for large ensembles reflect the underlying network connectivity. The localization matrix (6) is chosen, with $A_{ij} = |B_{ij}|$ since B has negative entries. We again use $\lambda = 0.46$ based on the coupling radius $r = 3$.

We find that localization again significantly improves the estimation capability of the EnKF, and yields accurate estimates for both the phases ϕ_i and the intrinsic parameters ζ_i when only 35 out of 50 phases are observed. This is shown in Fig. 4. Similar to Fig. 2(a), the RMS errors in phases ϕ_i (dashed) and intrinsic parameters ζ_i (solid) are shown in Fig. 4(a) over the course of the DA process. Results are shown for both the standard EnKF algorithm (blue) and the EnKF with localization (red). By time $t = 30$, localization has reduced the RMS errors in both the phases and the intrinsic parameters by more than a factor of 10. Similar to Fig. 2(b,c), the normalized residuals at the final time $t = 30$ are shown in Fig. 4(b,c) for the ϕ_i and ζ_i, respectively. The residuals clearly show the excellent performance of the EnKF approach with localization. Comparing the respective median residuals (horizontal lines in Fig. 4(b,c)), the localized algorithm yields more than a 100-fold decrease in error compared to the standard algorithm. As with the Kuramoto model, there is greater discrepancy for the unobserved nodes (open circles/triangles) compared to the observed nodes (filled circles/triangles), as expected. We note that for the quiescent, approximately stationary, nodes in the bump state (between $i = 30$ and $i = 40$), DA performs slightly worse. This is similar to the problem of synchronization discussed previously, though the persistent small fluctuations that occur in the theta neuron model allow better approximation than a purely stationary synchronized state.

Comparing the results for the theta neuron model (8) and the Kuramoto model (4), DA performs better, yielding smaller RMS errors, for the theta neuron model. This is because the dynamics of the theta neuron model are irregular, and, thus, the DA approach does not suffer from the data degeneracy issue that arises due to synchronization that plagues the Kuramoto model.

To summarize, data assimilation via the Ensemble Kalman Filter with network-specific localization and state space augmentation (and other minor modifications to account for periodicity of phase variables) can accurately determine both the phases and unknown model parameters in networks of coupled oscillators when only a subset of the phases are observed. Our novel localization approach utilizes the matrix exponential of the network's adjacency matrix, which encodes the connectivity between nodes. We have demonstrated the efficacy of our method for two types of oscillators; Kuramoto oscillators which approximate power grid dynamics, and theta neuron oscillators which approximate the dynamics in the brain. In both examples, data assimilation yields excellent approximations that closely agree with the truth.

Here we have taken random subsets for the observed nodes, but it is likely that there are optimal sets of nodes that should be observed. Real-world applications such as the power grid will also have physical limitations on which nodes can be observed. More work is needed to determine an optimal choice of observed nodes.

Here we have focused on using DA to learn the unknown intrinsic parameters for each oscillator (natural frequency or firing propensity). Future work should investigate whether DA can be used to determine unknown coupling functions or unknown network connectivity matrices. Unknown coupling functions can be estimated via Fourier series, with the unknown Fourier coefficients estimated via DA. Determining the network structure and the adjacency matrix is more challenging, since there is a greater number of unknowns, and the localization method presented here is dependent on an *a priori* known network structure. However, we believe that in cases

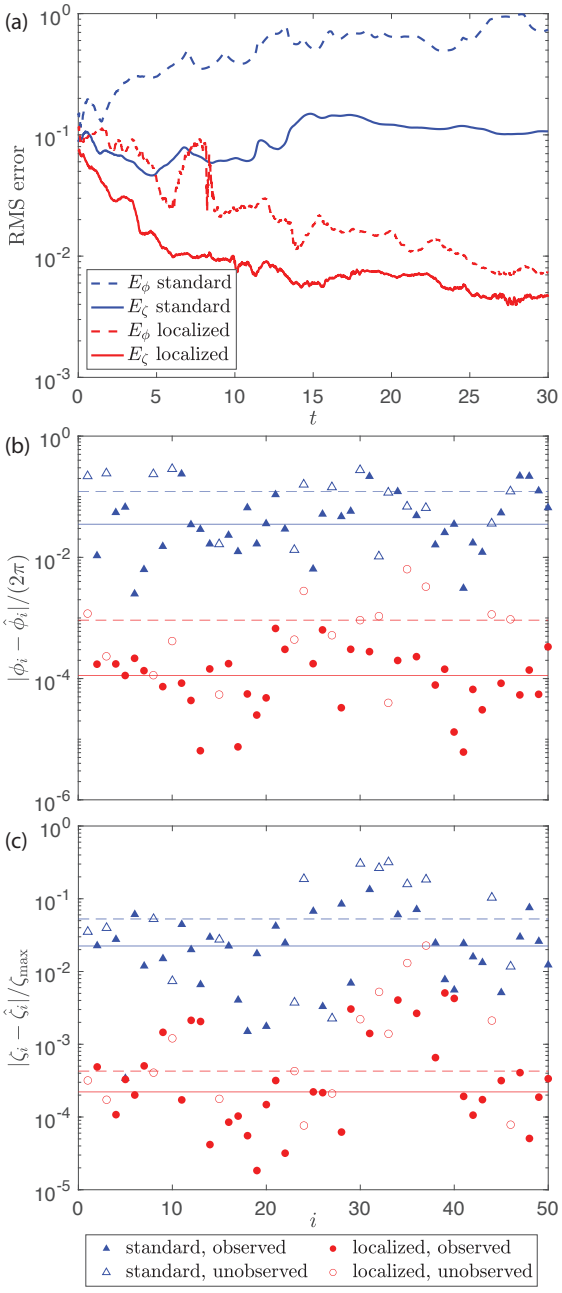

FIG. 4. (a) RMS errors (7) in ϕ (dashed) and ζ (solid) for DA applied to a theta neuron network (8) with 35 out of 50 phases observed. Results are shown for the standard EnKF with state space augmentation (blue) and EnKF with state space augmentation and network-based localization (6) with $\lambda = 0.46$ (red). (b), (c) Normalized residuals of ϕ_i and ζ_i, respectively, at final time $t = 30$ for the standard EnKF algorithm (blue triangles) and for the EnKF algorithm with localization (red circles). Filled symbols indicate observed nodes, and open symbols indicate unobserved nodes. Medians of the residuals are shown as horizontal lines; blue for standard, red for localized, solid for observed, dashed for unobserved. Note the logarithmic scale on the vertical axes. The theta neuron model parameters are $\kappa = 2$ and $\zeta_i \sim \mathcal{N}(-0.4, 0.1)$ with the coupling topology B described in the main text.

where the network topology is known reasonably well, e.g., it is known where power lines are in the power grid, but the non-zero weights of the adjacency matrix are not known, the localization can still be performed and the filter will converge. We remark that there has been significant progress in inferring the underlying network topology from incomplete observations[32–37].

Our novel localization method is specific to dynamics on networks, but not necessarily to coupled oscillators. We expect our localization approach to also improve DA when applied to other types of dynamics on networks, such as spreading processes (e.g., contagions)[38].

Appendix A: Ensemble Kalman filter

We describe here in detail the implementation of the Ensemble Kalman filter (EnKF). We recall that the states are labelled as $\boldsymbol{\phi} \in \mathbb{R}^N$ and the parameters are $\boldsymbol{\zeta} \in \mathbb{R}^N$ (here $\boldsymbol{\zeta}$ represents either the natural frequencies ω_i of the Kuramoto model or the intrinsic parameters ζ_i of the theta neuron model). To incorporate parameter estimation into an EnKF analysis step, an augmented state space $X = (\boldsymbol{\phi}^{\mathrm{T}}, \boldsymbol{\zeta}^{\mathrm{T}})^{\mathrm{T}} \in \mathbb{R}^{2N}$ is used.

In a Kalman filter, the analysis X_n^{a} at time t_n is the optimal estimate of the state given a forecast X_n^{f} and an observation Y_n. Treating X_n^{f} and X_n^{a}, $n \geq 0$ as random variables and assuming a Gaussian distribution for X_{n+1}^{f}, the analysis step for the mean $\overline{X}_n^{\mathrm{a}}$ is given by

$$\overline{X}_n^{\mathrm{a}} = \overline{X}_n^{\mathrm{f}} - K_n(H\overline{X}_n^{\mathrm{f}} - Y_n), \tag{9}$$

where the observation matrix $H = (H_\phi \ \mathbf{0}) \in \mathbb{R}^{N_{\mathrm{obs}} \times 2N}$ projects from (augmented) state space to the observation space. Here N_{obs} denotes the number of observed state variables, and $H_\phi \in \mathbb{R}^{N_{\mathrm{obs}} \times N}$ projects the phases onto observation space. If, for example, only the first phase ϕ_1 is observed, i.e $N_{\mathrm{obs}} = 1$, then $H = (1, 0, \cdots, 0) \in \mathbb{R}^{1 \times 2N}$. The Kalman gain matrix K_n is given by

$$K_n = P_n^{\mathrm{f}} H^T \left(H P_n^{\mathrm{f}} H^T + R\right)^{-1} \tag{10}$$

with forecast covariance matrix P_n^{f} and observational error covariance matrix R. We assume a diagonal observational error covariance with $R = \eta^2 I$.

We employ here a stochastic EnKF[11,13,39,40] to implement the Kalman analysis step. Ensemble Kalman filters allow for a dynamically adapted estimation of the forecast covariances, and they proved to be advantageous for nonlinear forward models, such as the Kuramoto model and the theta neuron model considered in this work, and for the non-Gaussian augmented state variables. Consider an ensemble of states $\boldsymbol{X} \in \mathbb{R}^{2N \times M}$ consisting of M members $X^{(i)} \in \mathbb{R}^{2N \times 1}$, $i = 1, \ldots, M$, that is,

$$\boldsymbol{X} = \left[X^{(1)}, X^{(2)}, \ldots, X^{(M)}\right], \tag{11}$$

with empirical mean

$$\overline{X} = \frac{1}{M}\sum_{i=1}^{M} X^{(i)}, \tag{12}$$

or, for phase variables, the mean given by (2), and associated matrix of ensemble deviations

$$\hat{\boldsymbol{X}} = \left[X^{(1)} - \overline{X}, X^{(2)} - \overline{X}, \ldots, X^{(M)} - \overline{X}\right]. \tag{13}$$

Ensembles for the forecast are denoted again by superscript f and those for the analysis by superscript a. In the forecast step the ϕ-component of each ensemble member is propagated independently using the non-linear forecast model (i.e., either the Kuramoto model or the theta-neuron model), updating the previous analysis ensemble $\boldsymbol{X}^{\mathrm{a}}_{n-1}$ to the next forecast ensemble $\boldsymbol{X}^{\mathrm{f}}_{n}$. The parameter component ζ of each ensemble member is considered constant during the forecast step, i.e. $\dot{\zeta} = 0$. The forecast covariance matrix P^{f}_n used in the analysis step (9) is estimated as a Monte-Carlo approximation from the forecast ensemble deviation matrix $\hat{\boldsymbol{X}}^{\mathrm{f}}_n$ via

$$P^{\mathrm{f}}_n = \frac{1}{M-1}\hat{\boldsymbol{X}}^{\mathrm{f}}_n \left(\hat{\boldsymbol{X}}^{\mathrm{f}}_n\right)^{\mathrm{T}} \in \mathbb{R}^{2N\times 2N}. \tag{14}$$

For periodic phase variables we apply the function $F(\theta) = \mathrm{mod}(\theta + \pi, 2\pi) - \pi$ to the deviations of phase variables as in (3).

To ensure that the analysis ensemble is statistically consistent with the Kalman filter, and, in particular, with the analysis error covariance $P^{\mathrm{a}} = (I - KH)P^{\mathrm{f}}$, in the stochastic ensemble Kalman filter[11,39] observations Y_n receive a stochastic perturbation $\boldsymbol{\eta}^{(i)}_n \in \mathbb{R}^{N_{\mathrm{obs}}\times 1}$, $i = 1, \ldots, M$, drawn independently from the Gaussian observational noise distribution $\mathcal{N}(\boldsymbol{0}, R)$. The associated ensemble of perturbed observations $\boldsymbol{Z}^{\mathrm{p}}_n \in \mathbb{R}^{N_{\mathrm{o}}\times M}$ is given by

$$\boldsymbol{Z}^{\mathrm{p}}_n = \left[Y_n - \boldsymbol{\eta}^{(1)}_n, Y_n - \boldsymbol{\eta}^{(2)}_n, \ldots, Y_n - \boldsymbol{\eta}^{(M)}_n\right]. \tag{15}$$

The EnKF analysis update step is then given by

$$\boldsymbol{X}^{\mathrm{a}}_n = \boldsymbol{X}^{\mathrm{f}}_n - K_n \Delta \boldsymbol{I}_n, \tag{16}$$

with the Kalman gain defined by (10) using (14) and the stochastic innovation

$$\Delta \boldsymbol{I}_n = H\boldsymbol{X}^{\mathrm{f}}_n - \boldsymbol{Z}^{\mathrm{p}}_n. \tag{17}$$

To mitigate against finite ensemble size effects we employ, as is typically done, covariance inflation with $P^{\mathrm{f}}_n \to \delta P^{\mathrm{f}}_n$ with $\delta = 1.001$[41]. This multiplicative inflation does not effect the ensemble mean but increases the forecast error covariance. In finite ensembles members may align, on short time scales, with the most unstable direction leading to an ensemble collapse in which many ensembles become indistinguishable. This leads to filter divergence in which the filter believes its own forecast since P^{f} is small, and the forecast is not corrected by new incoming observations. Apart from this ensemble collapse, finite ensembles typically generate spurious correlations. If two variables are not correlated, the Monte-Carlo approximation of the covariance yields entries of $\mathcal{O}(1/\sqrt{M})$. To mitigate against those spurious correlations which may spoil the analysis, localization is employed. Here, as outlined in the main text we use $P^{\mathrm{f}} \to \mathcal{L} \circ P^{\mathrm{f}}$ where $\circ$ denotes the Schur product and

$$\mathcal{L} = \begin{pmatrix} L & L \\ L & L \end{pmatrix} \tag{18}$$

is the localization matrix, where

$$L = \left(\mathrm{diag}(\mathcal{E})\right)^{-\frac{1}{2}} \mathcal{E} \left(\mathrm{diag}(\mathcal{E})\right)^{-\frac{1}{2}}, \tag{19}$$

and $\mathcal{E} = \exp(\lambda A)$, where λ is a tunable parameter and A is the network adjacency matrix, as in the main text (5)-(6). Since

$$\mathcal{E} = \exp(\lambda A) = \sum_{n=0}^{\infty} \frac{\lambda^n}{n!} A^n \tag{20}$$

is a weighted sum, with decreasing weights, of powers A^n, each counting the number of paths of length n, $\mathcal{E}$ encodes the network topology. Moreover, $\mathcal{E}$ is guaranteed to be a symmetric positive semi-definite matrix, i.e., a covariance matrix. The choice of λ will be discussed in Appendix B below.

The data assimilation cycles are initialized with an initial ensemble $\boldsymbol{X}^{\mathrm{a}}_0$ at $t = 0$ with $\boldsymbol{\phi}^{\mathrm{a}}_0 \sim \mathcal{N}(\boldsymbol{\phi}_0 + \boldsymbol{\eta}_\phi, \sigma^2_{\phi,1} I)$, such that $\boldsymbol{\phi}_0$ is the true value at $t = 0$ and $\boldsymbol{\eta}_\phi \sim \mathcal{N}(0, \sigma^2_{\phi,2}) \in \mathbb{R}^N$. In essence, $\sigma_{\phi,2}$ controls the perturbations of the ensemble means away from the truth, and $\sigma_{\phi,1}$ controls the spread of the ensemble around the ensemble means. The initial ensemble for the parameters $\boldsymbol{\zeta}$ has the analogous form $\boldsymbol{\zeta}^{\mathrm{a}}_0 \sim \mathcal{N}(\boldsymbol{\zeta}_0 + \boldsymbol{\eta}_\zeta, \sigma^2_{\zeta,1} I)$, such that $\boldsymbol{\eta}_\zeta \sim \mathcal{N}(0, \sigma^2_{\zeta,2}) \in \mathbb{R}^N$. For the Kuramoto model DA example in Fig. 2 we use $\sigma^2_{\phi,1} = \sigma^2_{\phi,2} = 0.25$, and $\sigma^2_{\omega,1} = \sigma^2_{\omega,2} = 0.25\,\sigma^2_\omega = 0.025$, where $\sigma^2_\omega = 0.1$ is the variance of the natural frequency distribution. For the theta neuron network example in Fig. 4 we use $\sigma^2_{\phi,1} = \sigma^2_{\phi,2} = 0.04$, and $\sigma^2_{\zeta,1} = \sigma^2_{\zeta,2} = 0.04\,\sigma^2_\zeta = 0.004$, where $\sigma^2_\zeta = 0.1$ is the variance of the intrinsic parameter distribution.

To test our DA methods we generate noisy observations Y_k by first simulating the model and then artificially adding noise with prescribed covariance $R = \eta^2 I$ with $\eta = 0.02$. To quantify the skill of the DA procedure, we compare in our numerical experiments the analysis mean with the true values.

Splitting of the analysis step into phase and parameter components

For our ensemble $\boldsymbol{X} = \left[X^{(1)}, X^{(2)}, \ldots, X^{(M)}\right]$, we can write each ensemble member in the form

$$X^{(j)} = \begin{pmatrix} \Phi^{(j)} \\ \Xi^{(j)} \end{pmatrix}, \tag{21}$$

such that $\Phi^{(j)} = \left(\phi_1^{(j)}, \ldots, \phi_N^{(j)}\right)^T$ and $\Xi^{(j)} = \left(\zeta_1^{(j)}, \ldots, \zeta_N^{(j)}\right)^T$. Analogous to (13), we can define ensemble deviations for the phases and parameters separately, i.e.,

$$\begin{aligned} \hat{\boldsymbol{\Phi}} &= \left[\Phi^{(1)} - \bar{\Phi}, \ldots, \Phi^{(M)} - \bar{\Phi}\right], \\ \hat{\boldsymbol{\Xi}} &= \left[\Xi^{(1)} - \bar{\Xi}, \ldots, \Xi^{(M)} - \bar{\Xi}\right], \end{aligned} \tag{22}$$

so that

$$\hat{\boldsymbol{X}} = \begin{pmatrix} \hat{\boldsymbol{\Phi}} \\ \hat{\boldsymbol{\Xi}} \end{pmatrix}. \tag{23}$$

From (14), it follows that the forecast covariance matrix is given by

$$P_n^{\mathrm{f}} = \frac{1}{M-1} \begin{pmatrix} \hat{\boldsymbol{\Phi}}_n^{\mathrm{f}} (\hat{\boldsymbol{\Phi}}_n^{\mathrm{f}})^T & \hat{\boldsymbol{\Phi}}_n^{\mathrm{f}} (\hat{\boldsymbol{\Xi}}_n^{\mathrm{f}})^T \\ \hat{\boldsymbol{\Xi}}_n^{\mathrm{f}} (\hat{\boldsymbol{\Phi}}_n^{\mathrm{f}})^T & \hat{\boldsymbol{\Xi}}_n^{\mathrm{f}} (\hat{\boldsymbol{\Xi}}_n^{\mathrm{f}})^T \end{pmatrix} = \begin{pmatrix} P_{\boldsymbol{\phi\phi}}^{\mathrm{f}} & P_{\boldsymbol{\phi\zeta}}^{\mathrm{f}} \\ P_{\boldsymbol{\zeta\phi}}^{\mathrm{f}} & P_{\boldsymbol{\zeta\zeta}}^{\mathrm{f}} \end{pmatrix}, \tag{24}$$

such that $P_{\boldsymbol{\phi\phi}}^{\mathrm{f}}$ is the covariance matrix of the phases only, $P_{\boldsymbol{\phi\zeta}}^{\mathrm{f}}$ is the covariance between the phases and parameters, etc.

Since we only observe the phases ϕ, we can separate the state and parameter update of the Kalman analysis step (9) as

$$\overline{\boldsymbol{\phi}}_n^{\mathrm{a}} = \overline{\boldsymbol{\phi}}_n^{\mathrm{f}} - P_{\boldsymbol{\phi\phi}}^{\mathrm{f}} H_{\boldsymbol{\phi}}^T \left(H_{\boldsymbol{\phi}} P_{\boldsymbol{\phi\phi}}^{\mathrm{f}} H_{\boldsymbol{\phi}}^T + R\right)^{-1} \Delta \boldsymbol{I}_n \tag{25a}$$

$$\overline{\boldsymbol{\zeta}}_n^{\mathrm{a}} = \overline{\boldsymbol{\zeta}}_n^{\mathrm{f}} - P_{\boldsymbol{\zeta\phi}}^{\mathrm{f}} H_{\boldsymbol{\phi}}^T \left(H_{\boldsymbol{\phi}} P_{\boldsymbol{\phi\phi}}^{\mathrm{f}} H_{\boldsymbol{\phi}}^T + R\right)^{-1} \Delta \boldsymbol{I}_n \tag{25b}$$

with innovation

$$\Delta \boldsymbol{I}_n := H_{\boldsymbol{\phi}} \overline{\boldsymbol{\phi}}_n^{\mathrm{f}} - Y_n. \tag{26}$$

We note that the covariance matrix $P_{\boldsymbol{\zeta\zeta}}^{\mathrm{f}}$, between the parameters and themselves, does not enter this equation.

Appendix B: Choosing λ for ring topology localization

The parameter λ in the localization matrix L (5) (cf. (19)-(20)) is important for controlling the rate of decay of correlations for distant nodes. The optimal value of λ could be found through numerical optimization, minimizing the RMS error of the DA algorithm across many simulations. Here we present a heuristic method to choose a value of λ for ring network topologies, and in Appendix C the method is generalized to random network topologies. Our heuristic is based on the typical decay of correlations that is observed from DA for large ensembles.

Considering a ring topology with connectivity radius r, we find that the correlations for large ensembles are negligible beyond a radius of $2r$. That is, for a node j, only the nodes k that have shortest path length from j to k less than or equal to 2 have significant correlations, and all other correlations can be considered spurious. This is shown in Fig. 5, which shows the average correlation as a function of distance from the node, scaled by the connectivity radius r. This data is obtained from the standard EnKF algorithm applied to the Kuramoto model with $N = 50$ and with a large ensemble $M = 100(2N + 1)$, as in Fig. 1(a) which shows correlations for $r = 3$. Fig. 5 shows that correlations become spurious beyond a distance of $2r$ ($d/r = 2$). Using this information, considering node $j = 1$, the nodes within a radius of $2r$, and which will have significant correlations are $k = 1, \ldots, 2r + 1$ and $k = N - r + 1, \ldots, N$, and all other nodes, i.e., $k = 2r + 2, \ldots, N - r$ will have spurious correlations. Knowing which correlations we wish to suppress, the next question is how much do we need to suppress them? For this, we choose an $0 < \epsilon < 1$, which will be the upper bound for the entries in L (5) beyond a radius of $2r$ from each node. Here we have chosen $\epsilon = 0.1$. Due to the symmetry of the ring topology, and the fact that $\exp(\lambda A)$ decays away from its diagonal, we require that

$$L_{1,2r+2} = \frac{\exp(\lambda A)_{1,2r+2}}{\exp(\lambda A)_{1,1}} \leq \epsilon, \tag{27}$$

where $\exp(\lambda A)_{i,j}$ denotes the i, j component of $\exp(\lambda A)$. Changing the inequality in (27) to an equality gives an equation that can be solved (numerically) for λ in terms r and ϵ. This is demonstrated schematically in Fig. 6, where the orange squares are the first spurious correlations, and have $L_{ij} = \epsilon$. In particular, in the first row we have $L_{1,2r+2} = \epsilon$.

For a fixed value of ϵ, we find that $1/\lambda$ scales approximately linearly with r, independent of N provided $r \ll N$. This is shown in Fig. 7. Fig. 7(a) shows that for fixed $\epsilon = 0.1$, the value of λ obtained from (27) does not depend on N, provided that $r \ll N$. Fig. 7(b) shows different approximately linear scalings between $1/\lambda$ and r for various values of ϵ (all using $N = 200$). This approximately linear scaling will be utilized to determine appropriate values of λ for general network topologies in Appendix C.

Appendix C: Data assimilation for random graphs

In this section we show results for DA applied to the Kuramoto model with random network topologies. First we discuss the localization procedure, then we show that localization greatly improves estimation accuracy.

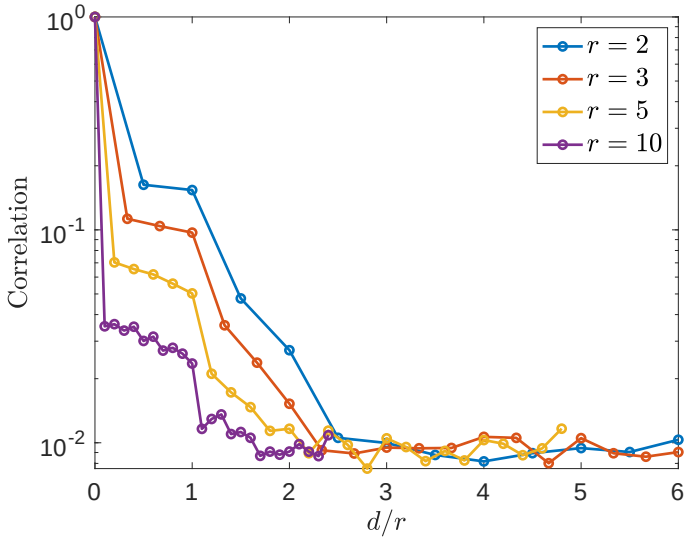

FIG. 5. Average forecast correlation (averaging P_k^{f} over all nodes and all time steps) as a function of distance d from the node, scaled by the coupling radius r, for $r = 2, 3, 5, 10$. All results are obtained from runs of the standard EnKF algorithm applied to the Kuramoto model ($N = 50$, $\kappa = 80/r$, $\omega_i \sim \mathcal{N}(0, 0.1)$ in a ring topology), with a large ensemble $M = 100(2N + 1)$. Note the coupling strength $\kappa = 80/r$ is used to maintain a constant effective coupling strength $\kappa/\langle d \rangle$, where $\langle d \rangle$ is the mean degree.

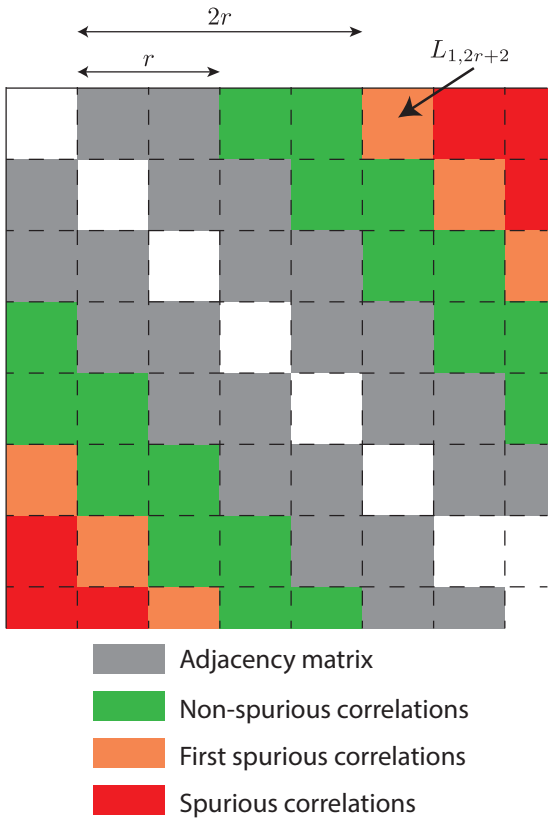

FIG. 6. Schematic diagram for the condition (27) that defines λ. Shown is the top right corner of the matrix L, the gray squares indicate the underlying adjacency matrix (here a ring with $r = 2$), the green squares indicate correlations that should not be considered spurious because they are within a distance of $2r$ from the node, and orange/red squares indicate correlations that should be considered spurious because the distance from the node is greater than $2r$. The orange squares are those whose value will be ϵ for λ satisfying (27) with equality, e.g., in the first row $L_{1,2r+2} = \epsilon$.

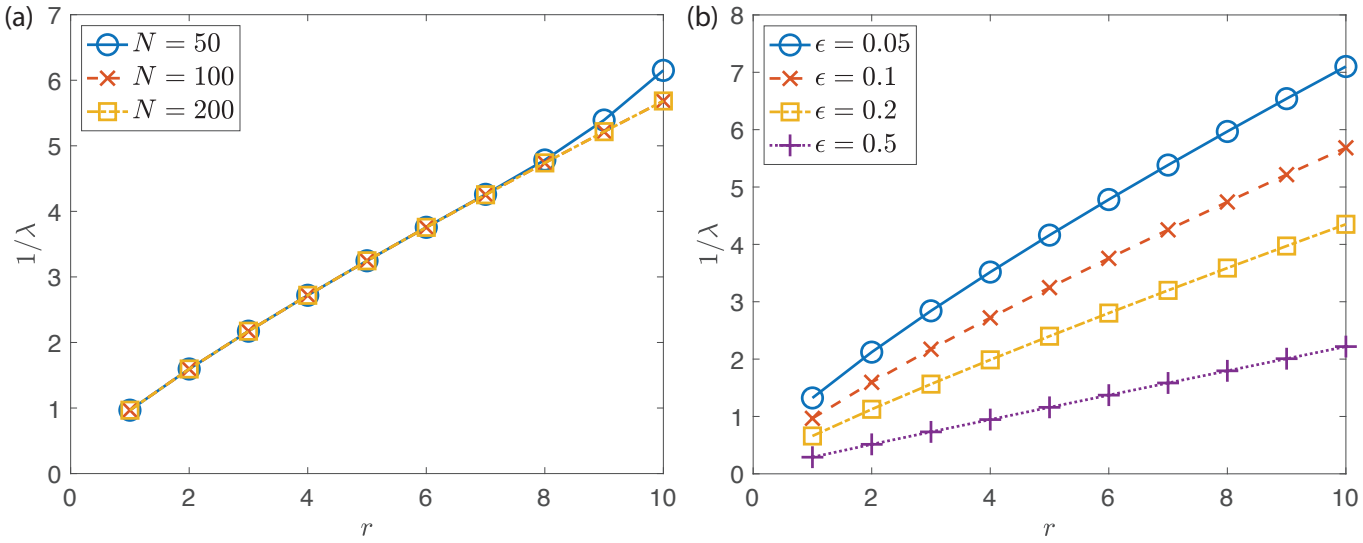

FIG. 7. $1/\lambda$ as a function of r for various values of N and ϵ. In all cases λ is obtained from (27). (a) Fixed $\epsilon = 0.1$, $N = 50, 100, 200$. (b) Fixed $N = 200$, $\epsilon = 0.05, 0.1, 0.2, 0.5$.

Localization method for random graphs

In the main text it was shown that localization using (6) greatly improves the estimation accuracy of the EnKF algorithm for the Kuramoto model (4) with a ring network topology. Here we show that localization also yields a large improvement for random graphs. We focus on Erdős-Renyi (ER) graphs, for which each pair of nodes are coupled with probability p (here we us $p = 0.1$).

We determine the value of λ to use in the localization matrix (5) based on a ring topology with the equivalent mean degree. For an ER network with N nodes, the mean degree is $(N-1)p$. The corresponding ring topology with the same mean degree has connectivity radius $r^* = (N-1)p/2$ (ignoring for now that this is likely not an integer). We now find λ^* corresponding to the ring topology with $r = r^*$. If r^* is an integer, we follow the procedure from Appendix B, otherwise, we interpolate to find λ^*. We could linearly interpolate between the two nearest integer values of r, but we can do better by noting that $1/\lambda$ scales approximately linearly with r, as shown in Fig. 7. Thus, we find the fit of the form

$$\frac{1}{\lambda} = mr + c, \tag{28}$$

passing through $(r_1, 1/\lambda_1)$ and $(r_2, 1/\lambda_2)$, where $r_1 = \lfloor r^* \rfloor$ and $r_2 = \lceil r^* \rceil$ are the two nearest integers to r^*, and λ_1 and λ_2 are the corresponding values of λ obtained from Appendix B That is,

$$m = \frac{1/\lambda_2 - 1/\lambda_1}{r_2 - r_1}, \quad c = \frac{1}{\lambda_1} - mr_1. \tag{29}$$

The value λ^* is then given by

$$\lambda^* = (mr^* + c)^{-1}. \tag{30}$$

As an example, for an ER network with $N = 50$ and $p = 0.1$, the mean degree is 4.5. The equivalent ring topology has $r^* = 2.45$. The two nearest integers are $r_1 = 2$ and $r_2 = 3$. From the method in Appendix B we obtain $\lambda_1 = 0.627$ and $\lambda_2 = 0.460$. Substituting these values into (29)-(30) yields $\lambda^* = 0.539$.

The effect of localization for Erdős-Renyi networks

To show the improvement that is gained by employing localization, we consider 500 realizations of the Kuramoto model with a random network topology (ER graph with $N = 50$, $p = 0.1$) and natural frequencies ω_i (drawn from $\mathcal{N}(0, 0.1)$). For each realization we run the EnKF algorithm with 35 out of 50 nodes observed, both with and without localization. We record the RMS errors in both the phases ϕ_i and the frequencies ω_i at the end time $t = 15$ (which is the typical synchronization time for such a network with coupling strength $\kappa = 10$). The results are shown in Fig. 8, which shows histograms of RMS errors in (a) ϕ, and (b) ω. As for the ring topology, we see a significant improvement in the estimation accuracy when localization is employed.

Appendix D: Comparison between matrix exponential localization and Gaspari-Cohn localization functions

Localization has been widely employed in applications of the EnKF. However, most applications assume an underlying Euclidean geometry such that the distance between nodes is their Euclidean distance. Based on the assumption of Euclidean geometry, methods have been created to generate localization functions. One method is that of Gaspari & Cohn[29], which yields the commonly used localization function

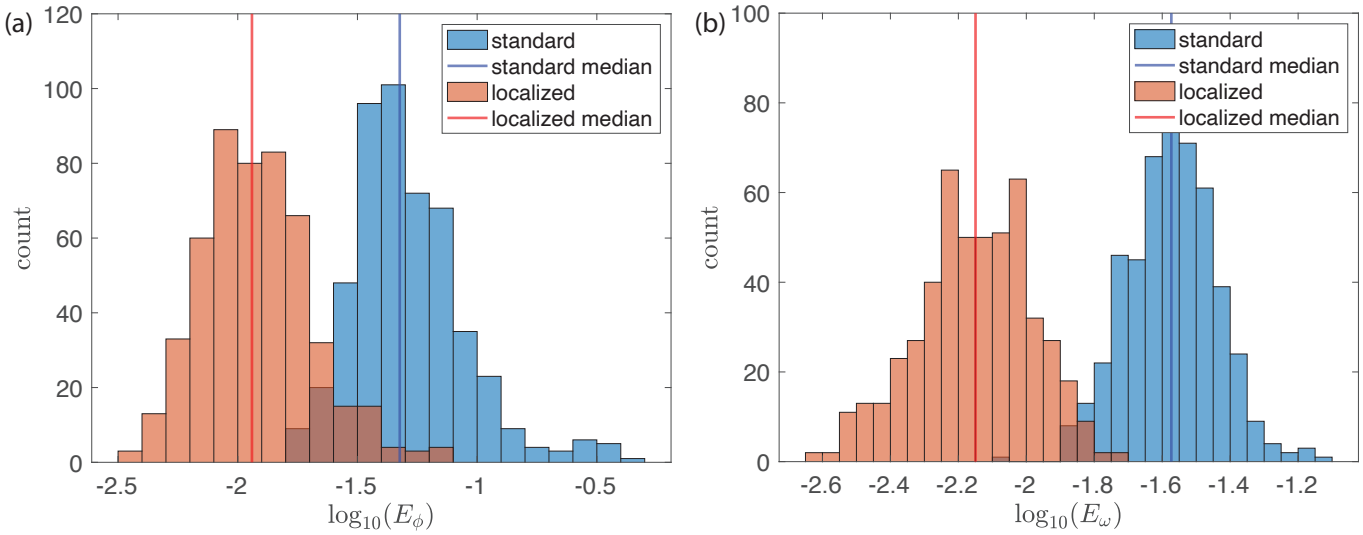

FIG. 8. RMS errors at $t = 15$ in (a) ϕ, and (b) ω for 500 realizations of DA applied to the Kuramoto model with 35 out of 50 phases observed, and an Erdős-Renyi network topology ($N = 50$ and $p = 0.1$). Histograms are shown for the standard EnKF algorithm (blue) and the EnKF algorithm with network-specific localization (5) with $\lambda = \lambda^* = 0.539$ obtained from (30). For both algorithms the ensemble size used is $M = 101$. Also shown are the respective median values for each distribution – blue vertical lines show medians for the standard algorithm, and red vertical lines show medians for the localized algorithm. The coupling strength for all simulations is $\kappa = 10$ and the natural frequencies are drawn from $\mathcal{N}(0, 0.1)$.

$$C(z,c) = \begin{cases} -\frac{1}{4}\left(\frac{|z|}{c}\right)^5 + \frac{1}{2}\left(\frac{z}{c}\right)^4 + \frac{5}{8}\left(\frac{|z|}{c}\right)^3 - \frac{5}{3}\left(\frac{z}{c}\right)^2 + 1, & 0 \le |z| \le c, \\ \frac{1}{12}\left(\frac{|z|}{c}\right)^5 - \frac{1}{2}\left(\frac{z}{c}\right)^4 + \frac{5}{8}\left(\frac{|z|}{c}\right)^3 + \frac{5}{3}\left(\frac{z}{c}\right)^2 - 5\left(\frac{|z|}{c}\right) + 4 - \frac{2}{3}\left(\frac{|z|}{c}\right)^{-1}, & c \le |z| \le 2c, \\ 0, & 2c \le |z|, \end{cases} \tag{31}$$

where $|z|$ is the Euclidean distance between the nodes and c is a length-scale that controls the decay of correlation. The function C is a 5th-order piecewise rational function with compact support. In generating a localization matrix L, one would define $L_{ij} = C(|x_i - x_j|, c)$, where x_i is the position of node i in space, and $|x|$ denotes the Euclidean norm. Note, the proof that C is a correlation function relies on the Euclidean distance being used.

When considering localization on networks, there are many possible definitions of the "distance" between two nodes. Possibly the simplest definition is the shortest path length, that is, the distance $d(i, j)$ between nodes i and j is the length of the shortest path connecting i to j. With this definition, one can, in principle, define the function C from (31), and create a matrix $\tilde{L}_{ij} = C(d(i, j), c)$, where here c is a characteristic length for meaningful correlations. However, this construction does not in general yield a correlation matrix, in particular, the matrix $\tilde{L}$ is not always positive semi-definite. Nevertheless, we can compare the matrix $\tilde{L}$ with L obtained from (5). In the previous section we noted that for ring topologies correlations appear to be spurious for node distances greater than 2, so a choice $c \approx 2$ is logical. Fig. 9 shows the first rows of the respective localization matrices L and $\tilde{L}$ for the ring topology with $N = 50$ and $r = 3$ using the matrix exponential construction (5) (black crosses) compared to the Gaspari-Cohn construction with shortest path length distance (red circles). There is very good agreement between the two constructions. We note that the matrix exponential has a 'smoother' decay, which more accurately reflects the decay of correlations observed in Fig. 5.

The similarity between the matrix exponential and the Gaspari-Cohn function with shortest path length distance (31) suggests that the matrix exponential is a good choice for localization. One advantage of the Gaspari-Cohn construction (31) is that the function is compactly supported, resulting in a sparse localization matrix, which in turn reduces the computational cost of data assimilation. One could produce a sparse matrix from the localization matrix L derived from the matrix exponential (5) by setting all entries below a threshold value equal to zero, however, the resulting matrix is not guaranteed to be positive semi-definite. Alternatively, a sparse localization matrix can be obtained by considering

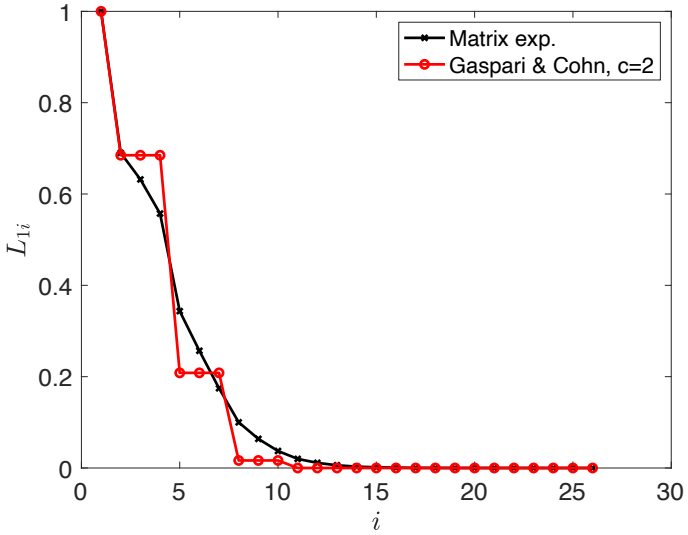

FIG. 9. Values L_{1i} from the first row of the matrix exponential correlation matrix shown in Fig. 1(b) using (5) (cf. (19)) with $\lambda = 0.46$ (black crosses) and values $\tilde{L}_{1i} = C(d(1,i),c)$ using the Gaspari-Cohn function with shortest path length distance (31) with $c = 2$ (red circles). Both use an underlying ring topology with $N = 50$ and $r = 3$.

a normalization of a matrix of the form

$$B = \sum_{i=0}^{n} \alpha_i A^i \tag{32}$$

which is positive definite provided that A is symmetric and $\sum_{i=0}^{n} \alpha_n \eta^i \geq 0$ for all eigenvalues η of A, and is sparse provided that $n \ll N$. This construction allows for $n+1$ tunable parameters α_i. For example, using the matrix exponential as a guide, one could consider the truncation

$$B = \sum_{i=0}^{n} \frac{\lambda^i}{i!} A^i \tag{33}$$

for $n \ll N$ and with one parameters λ. However, one would have to check that $\sum_{i=0}^{n} \frac{\lambda^i}{i!} \eta^i$ is non-negative for all eigenvalues η of A. Consideration of sparse localization matrices of this form is left as an avenue for future work.

Appendix E: Varying the number of observed nodes

In the main text we considered a ring network topology such that the observed nodes were chosen at random, with 35 out of 50 phases being observed. Here we again use a ring topology, and we investigate the effect of changing the number of observed phases when DA is applied to the Kuramoto model. To simplify the analysis, we do not choose observed nodes at random, but rather require that they are evenly distributed, e.g. every fourth node in the ring is observed/unobserved. Note, there is still inhomogeneity in the natural frequencies ω_i.

We consider here $N = 60$ which allows more data points for evenly distributed observed/unobserved nodes. Note that all natural frequencies remain unobserved and are initially unknown, as is assumed in the main text. RMS errors in ϕ (dashed) and ω (solid) are shown in Fig. 10 for different fractions of observed phases, for both the standard EnKF algorithm (blue) and the EnKF algorithm with network-specific localization (red). Both DA methods use $M = 2N + 1 = 121$ ensemble members. As for the results presented in the main text, we see a significant decrease in the RMS error when localization is employed, and the improvement in accuracy becomes most pronounced when at least half of the phases are observed. The standard EnKF algorithm shows little reduction in RMS errors until the fraction of observed nodes is very close to 1, whereas the localized algorithm shows an improving trend starting when approximately one third of the phases are observed.

The observed non-monotonic trend in the RMS error is likely due to the inhomogeneity in the natural frequencies.

* lauren.smith@auckland.ac.nz
† georg.gottwald@sydney.edu.au

[1] C. Bick, M. Goodfellow, C. R. Laing, and E. A. Martens, Understanding the dynamics of biological and neural oscillator networks through exact mean-field reductions: a review, The Journal of Mathematical Neuroscience **10**, 1 (2020).
[2] E. Montbrió, D. Pazó, and A. Roxin, Macroscopic description for networks of spiking neurons, Phys. Rev. X **5**, 021028 (2015).
[3] H. Schmidt and D. Avitabile, Bumps and oscillons in networks of spiking neurons, Chaos: An Interdisciplinary

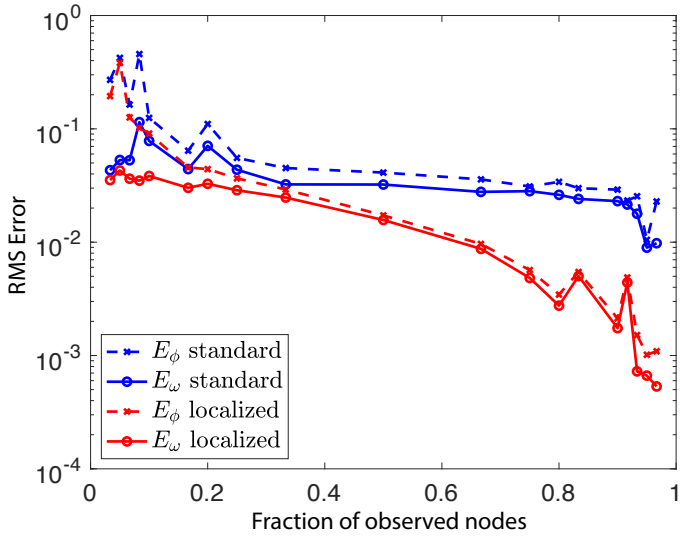

FIG. 10. RMS errors at $t = 15$ in ϕ (dashed) and ω (solid) for different fractions of observed phases when the EnKF is applied to the Kuramoto model with $N = 60$, $\kappa = 10$, $\omega_i \sim \mathcal{N}(0, 0.1)$ and a ring network topology with $r = 3$. Results for the standard EnKF algorithm are shown in blue, results for the EnKF algorithm with network-specific localization are shown in red. Both DA approaches use $M = 2N + 1 = 121$ ensemble members.

Journal of Nonlinear Science **30**, 033133 (2020).

[4] J. Machowski, J. W. Bialek, and J. Bumby, *Power system dynamics: stability and control* (John Wiley & Sons, 2011).

[5] T. Nishikawa and A. E. Motter, Comparative analysis of existing models for power-grid synchronization, New J. Phys. **17**, 10.1088/1367-2630/17/1/015012 (2015).

[6] G. Filatrella, A. H. Nielsen, and N. F. Pedersen, Analysis of a power grid using a Kuramoto-like model, Eur. Phys. J. B **61** (2008).

[7] B. Schäfer and G. C. Yalcin, Dynamical modeling of cascading failures in the Turkish power grid, Chaos **29**, 10.1063/1.5110974 (2019).

[8] G. Evensen, Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics, J. Geophys. Res.-Oceans **99** (1994).

[9] P. L. Houtekamer and H. L. Mitchell, Data assimilation using an ensemble Kalman filter technique, Mon. Weather Rev. **126** (1998).

[10] A. H. Jazwinski, *Stochastic processes and filtering theory* (Courier Corporation, 2007).

[11] G. Evensen, *Data assimilation: The ensemble Kalman filter*, Vol. 2 (Springer, 2009).

[12] G. Evensen, The ensemble Kalman filter for combined state and parameter estimation, IEEE Contr. Syst. Mag. **29** (2009).

[13] S. Reich and C. Cotter, *Probabilistic forecasting and Bayesian data assimilation* (Cambridge University Press, 2015).

[14] K. Law, A. Stuart, and K. Zygalakis, *Data assimilation: A mathematical introduction* (Springer-Verlag, New York, 2015).

[15] J. Carrillo, F. Hoffmann, A. Stuart, and U. Vaes, The ensemble Kalman filter in the near-Gaussian setting, arXiv preprint arXiv:2212.13239 (2022).

[16] M. J. Panaggio, M.-V. Ciocanel, L. Lazarus, C. M. Topaz, and B. Xu, Model reconstruction from temporal data for coupled oscillator networks, Chaos **29**, 10.1063/1.5120784 (2019).

[17] S. G. Shandilya and M. Timme, Inferring network topology from complex dynamics, New J. Phys. **13**, 10.1088/1367-2630/13/1/013004 (2011).

[18] A. Pikovsky, Reconstruction of a random phase dynamics network from observations, Phys. Lett. A **382** (2018).

[19] K. Owens and J. N. Kutz, Data-driven discovery of governing equations for coarse-grained heterogeneous network dynamics, arXiv preprint arXiv:2205.10965 (2022).

[20] Y. Kuramoto, *Chemical Oscillations, Waves, and Turbulence* (Springer-Verlag, Berlin, 1984).

[21] S. H. Strogatz, From Kuramoto to Crawford: Exploring the onset of synchronization in populations of coupled oscillators, Physica D **143**, 10.1016/S0167-2789(00)00094-4 (2000).

[22] C. R. Laing, Derivation of a neural field model from a network of theta neurons, Phys. Rev. E **90**, 010901 (2014).

[23] C. R. Laing, The dynamics of networks of identical theta neurons, The Journal of Mathematical Neuroscience **8**, 1 (2018).

[24] C. R. Laing and O. Omel'chenko, Moving bumps in theta neuron networks, Chaos: An Interdisciplinary Journal of Nonlinear Science **30**, 043117 (2020).

[25] O. Omel'chenko and C. R. Laing, Collective states in a ring network of theta neurons, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences **478**, 20210817 (2022).

[26] J. L. Anderson, An ensemble adjustment Kalman filter for data assimilation, Monthly Weather Review **129**, 2884 (2001).

[27] P. L. Houtekamer and H. L. Mitchell, A sequential ensemble Kalman filter for atmospheric data assimilation, Monthly Weather Review **129**, 123 (2001).

[28] F. A. Rodrigues, T. K. D. Peron, P. Ji, and J. Kurths, The Kuramoto model in complex networks, Phys. Rep.

610 (2016).
[29] G. Gaspari and S. E. Cohn, Construction of correlation functions in two and three dimensions, Quarterly Journal of the Royal Meteorological Society **125**, 723 (1999).
[30] P. C. Bressloff, Spatiotemporal dynamics of continuum neural fields, Journal of Physics A: Mathematical and Theoretical **45**, 033001 (2012).
[31] K. Wimmer, D. Q. Nykamp, C. Constantinidis, and A. Compte, Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory, Nature Neuroscience **17**, 431 (2014).
[32] T. P. Peixoto, Reconstructing networks with unknown and heterogeneous errors, Phys. Rev. X **8**, 041011 (2018).
[33] T. P. Peixoto, Network reconstruction and community detection from dynamics, Phys. Rev. Lett. **123**, 128301 (2019).
[34] M. Tyloo, R. Delabays, and P. Jacquod, Reconstructing network structures from partial measurements, Chaos: An Interdisciplinary Journal of Nonlinear Science **31**, 103117 (2021).
[35] W. Deng, C. Yang, K. Huang, and W. Wu, A two-stage reconstruction method for complex networked system with hidden nodes, Chaos: An Interdisciplinary Journal of Nonlinear Science **32**, 053105 (2022).
[36] M. Chen, Y. Zhang, Z. Zhang, L. Du, S. Wang, and J. Zhang, Inferring network structure with unobservable nodes from time series data, Chaos: An Interdisciplinary Journal of Nonlinear Science **32**, 013126 (2022).
[37] T. Gaskin, G. A. Pavliotis, and M. Girolami, Inferring networks from time series: a neural approach, arXiv preprint arXiv:2303.18059 (2023).
[38] M. De Domenico, C. Granell, M. A. Porter, and A. Arenas, The physics of spreading processes in multilayer networks, Nature Physics **12**, 901 (2016).
[39] G. Burgers, P. J. van Leeuwen, and G. Evensen, Analysis scheme in the ensemble Kalman filter, Monthly Weather Review **126**, 1719 (1998).
[40] K. Law, A. Stuart, and K. Zygalakis, *Data assimilation: A mathematical introduction* (Springer-Verlag, New York, 2015).
[41] J. L. Anderson and S. L. Anderson, A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts, Monthly Weather Review **127**, 2741 (1999).