

Response to the Reviewers

A Continuous-Time Ensemble Kalman–Bucy Smoother for Causal Inference and Model Discovery

Dear Editor and Reviewers,

Thank you for your careful reading of our manuscript and for the thoughtful and constructive feedback! We have carefully considered each comment and revised the manuscript accordingly. Below, we respond to each point individually and indicate the corresponding changes and their location in the revised manuscript. We also provide a separate comparison PDF highlighting the changes made relative to the original submission.

We are grateful for your time and care devoted to the review. Your comments and recommendations have helped us improve the clarity, precision, and presentation of the manuscript.

Here, quoted equation, figure, table, and section numbers refer to the original manuscript, whereas line numbers given in our responses refer to the revised manuscript.

Response to Reviewer 1

This paper introduces a new data assimilation method called EnKBS, which improves the ability to estimate the hidden states of complex systems by using information from both the past and the future instead of relying only on current observations. The authors argue that traditional filtering methods often struggle when systems change rapidly or experience sudden events since they can only react after those changes become visible. Their proposed smoother also works in continuous time, avoids expensive derivative calculations. Overall I found the proposed idea novel and the contribution is of clear value, here lists my questions/comments before I can recommend the publication of this manuscript.

We thank Reviewer 1 for their careful reading of the manuscript and for the constructive and encouraging comments. We especially appreciate their positive assessment of the novelty and value of the proposed approach. Their questions and suggestions have helped us sharpen the discussion of the applicability, robustness, and possible extensions of the EnKBS framework.

Comment 1: Non-Gaussian distributions

Since EnKBS mainly relies on ensemble means and covariances, it is unclear to me whether it can accurately capture non-Gaussian distribution. Could the authors add some discussion in the manuscript?

Response:

Thank you for raising this important concern. In general, the EnKBS is not guaranteed to fully represent arbitrary non-Gaussian distributions. Although the ensemble itself is not constrained to be Gaussian and may retain some non-Gaussian features, the EnKBS update is based on empirical

first- and second-order moments and should be viewed as a moment-based approximation, like other ensemble Kalman methods.

This is also reflected in the dyad experiment: in the $u \rightarrow v$ direction, the non-injective term $-cu^2$ can cause ensemble analysis members to split into opposite-sign branches, increasing the variance and estimation error. We have added a brief discussion in the manuscript at lines 124–127 to make this limitation explicit.

Comment 2: Sensitivity to inflation and localization

Since inflation/localization techniques are crucial for stability, it would be useful to know whether the method remains reliable when inflation/localization parameters are chosen differently.

Response:

We appreciate this question about parameter sensitivity. Tables 1 and 2 already provide a systematic sensitivity analysis by varying both the localization radius and the inflation factor in the Lorenz-96 experiment. Over a broad range of moderate parameter choices, the EnKBS remains stable and yields comparable RMSEs, indicating that the reported performance does not depend on a narrowly tuned configuration. Localization has a more pronounced effect on EnKBS stability, whereas mild inflation mainly provides a modest improvement in RMSE. More aggressive choices, such as excessive inflation or a very large localization radius, can lead to numerical divergence.

The role of these parameters is also problem dependent. In the low-dimensional dyad example, localization is unnecessary, while inflation has no noticeable effect in the regime considered and is therefore not used.

Comment 3: Coupling with AI-enhanced data assimilation

It would be beneficial if the authors could discuss if/how the proposed EnKBS could be coupled with modern AI-enhanced data assimilation schemes.

Response:

This is an excellent suggestion—thank you! One natural interface with existing ML-enhanced DA schemes is through the forecast model. In applications where repeated integration of the dynamical model is computationally expensive, a learned emulator or data-driven surrogate could replace or augment the physical forecast model while retaining the EnKBF/EnKBS assimilation framework (Hodyss and Morzfeld, 2026).

A second possibility is to learn a surrogate for the backward correction in the EnKBS itself. In the current formulation, the smoother involves the covariance-based correction $\Sigma P_f^{-1}(x_s^{(i)} - x_f^{(i)})$, which requires access to the filtering ensemble trajectory and repeated evaluation of empirical covariance information. A learned surrogate for this term could potentially reduce the memory and computational cost of the backward pass. More generally, such a learned correction need not be restricted to the present linear, moment-based form: a nonlinear and state-dependent correction could potentially exploit information beyond the first two moments and provide greater flexibility for non-Gaussian distributions. Related score-based, ML-enhanced DA methods are also being

explored in ongoing work by one of the authors, with a manuscript in preparation. We have added a brief discussion of these possibilities to the Conclusion, lines 457–462.

Comment 4: Dependence on ensemble size

How dependent is performance on ensemble size (e.g., compared to traditional EnKF approaches)? The paper highlights success with ensembles of $O(10)$, but under what conditions does performance sharply deteriorate?

Response:

Thank you for raising this point. The ensemble-size dependence is examined in all three numerical examples and the results suggest that there is no universal critical ensemble size below which performance deteriorates sharply. Rather, the behavior depends on the system dimension, nonlinearity, and regularization.

In the Lorenz–96 experiment of Section 3.1, the EnKBS is more sensitive to small ensembles than the corresponding EnKBF. Without localization, reducing the ensemble size can lead to instability of the backward smoother even when the forward filter remains stable, whereas localization allows stable smoothing with much smaller ensembles compared to unlocalized EnKBS (down to $m = 5$ in the tested setting). Thus, the most pronounced deterioration occurs in the high-dimensional case when the ensemble is too small to provide a sufficiently reliable covariance estimate for the backward update. For the dyad model of Section 3.2, Figure 4 provides a direct ensemble-size convergence study against the analytic CGNS filter and smoother. As m increases, the EnKBS statistics and the associated ACI diagnostics converge toward the analytic benchmark. Even with $m = 10$, the main qualitative ACI/CIR patterns are retained. In the Lorenz–84 model learning experiment of Section 3.3, $m = 10, 50$, and 100 produce similar learning behavior, with the $m = 10$ case typically converging a few iterations later.

These ensemble-size effects are already discussed in the corresponding numerical sections of the manuscript.

Minor Comment: Terminology and presentation

In the paper the authors alternate between “DA methods”, “filtering methods”, and “smoothing algorithms”. It would be better to maintain more consistent terminology.

The paper is overall well written but please check some minor issues, e.g., “allows for broad application(s)”, “mean and covariance equation(s)”, etc.

Response:

Thank you for the careful reading. We have revised the terminology throughout the manuscript for consistency and corrected the noted grammatical issues, including “broad applications” and “mean and covariance equations”, together with other minor wording issues.

Response to Reviewer 2

The authors propose continuous-time and derivative-free schemes, namely ensemble Kalman-Bucy filter (EnKBF) / smoother (EnKBS), in which the conditional (smoothing) distributions are reconstructed from ensemble moments, avoiding tangent-linear and adjoint models. The smoother is claimed to converge to the exact smoothing solution in linear dynamic systems. The proposed approaches have employed covariance localization and inflation to address high-dimensional problems. In this manuscript, the authors illustrated the performance of EnKBF/EnKBS on state estimation for high-dimensional chaotic systems (Lorenz 96, Lorenz 84) and iterative learning for causal relationships and their temporal influence range in a dyadic trigger-feedback model and in a reduced-order model of midlatitude atmospheric circulation. The topic suits the NPG scope as continuous-time smoothers for nonlinear systems are likely under-explored, and the link between smoothing and backward tracing of causal drivers to extreme events is an appealing idea. However, several methodological points in the current manuscript could be elaborated, and the performance of the proposed approaches relative to existing ensemble smoothers should be explicitly demonstrated. I recommend addressing the following points before further consideration.

We thank Reviewer 2 for the thorough and thoughtful review of our manuscript. We particularly appreciate your recognition that continuous-time smoothers for nonlinear systems remain comparatively under-explored and that the connection between smoothing and the backward tracing of causal drivers is an appealing aspect of the framework. The methodological questions raised in the review are very helpful and have encouraged us to clarify several important points, especially the interpretation of the continuous-time formulation, the backward dynamics, and the precise scope of the paper’s contribution.

Major Comment 1: Nonlinearity, non-Gaussianity, and relation to existing smoothers

In the abstract, the authors state that integrating future observations reveals the causal mechanism and supports retrospective updates. The Kalman-Bucy smoother in the linear-Gaussian case already provides this benefit. For the dyadic trigger-feedback model and the reduced-order atmospheric model, please clarify:

- (i) whether the dynamics are genuinely nonlinear and non-Gaussian in the regimes considered, and*
- (ii) what the EnKBS provides that the analytical KS (or a linearized/unscented RTS smoother) does not.*

Response:

Thank you for drawing attention to these points. We address these separately:

- (i) All three dynamical models considered in this work are genuinely nonlinear and exhibit non-Gaussian statistics. Specifically, for the dyad model of Section 3.2, the intermittent extreme events generated by the trigger-feedback mechanism lead to non-Gaussianity. The empirical probability density functions (PDFs) are shown in Figure 1, where the non-Gaussianity is mostly visible in the skewness and asymmetry of u . The Lorenz-84 system of Section 3.3 is also genuinely non-Gaussian. Its non-Gaussian behavior is visible in the empirical PDFs already shown in Figure 9, showing the multimodal distributions of x , y and z .




Figure 1: Trajectories and empirical PDFs of the dyad variables u and v . The trajectories are generated using the same model parameters and numerical configuration as in the dyad experiments in the manuscript. The empirical PDFs are estimated from a trajectory of length 10000 time units, while only a segment of that trajectory is displayed. The dashed black curves denote Gaussian distributions constructed using the empirical means and variances.

- (ii) Relative to linearized or extended RTS smoothers, the EnKBS avoids explicit tangent-linear or adjoint models of the nonlinear dynamics. Unscented RTS smoothers also avoid such linearization by propagating sigma points. However, both extended and unscented RTS approaches are primarily formulated to propagate the smoothing mean and covariance. In the unscented case, the sigma points are propagated and regenerated to approximate the corresponding moments. In contrast, the EnKBS evolves a stochastic analysis ensemble, providing trajectory-level samples for an approximate but comprehensive distributional description of the smoother posterior rather than only its first two moments at each time. This distinction is particularly important for the model-discovery application, where the individual smoothing trajectories are used as inputs to the iterative learning procedure. Replacing these trajectories by the smoother mean alone would remove the realization-to-realization variability and underestimate the uncertainty and temporal correlations/memory associated with the hidden stochastic dynamics.

Major Comment 2: Definition of the assimilation pseudo-time s

The introduction of assimilation time s in pages 5–6 needs a proper definition. As written, s is called a "time interval" although it appears to be a pseudo-time coordinate. Please clarify: does s depend on t ? How exactly does s enter Eqs. (2) and (3), which currently show no explicit s -dependence? Is P in these equations the propagated filtering covariance, and if so, is it a function of t , of s , or of both?

Response:

Thank you for pointing out this ambiguity. The original motivation for introducing the pseudo-time

s was to interpret the analysis step as a homotopy: At each fixed physical time t , an ODE flow in $s \in [0, 1]$ continuously maps the forecast ensemble at $s = 0$ to the analyzed ensemble at $s = 1$.

However, this pseudo-time construction served only as a formal interpretation of the EnKBF. Since this is not essential in the subsequent EnKBS development and might lead to ambiguity in the notations, we decided to remove it from Section 2.2. The manuscript now introduces the stochastic EnKBF directly in physical time, and the corresponding discrete-to-continuous derivation in Appendix A has also been revised accordingly.

Major Comment 3: Backward stochastic dynamics, discretization, and stability

Are backward dynamics reversible in time as always? Equation (9) uses reverse prediction of $x(t)$ given $x(t+1)$. Simply integrating the forward drift f backward in time is not, in general, the time reversal of a stochastic process: the reverse-time SDE has a modified drift and the forward noise process is not reversible as such [Anderson (1982)]. The authors should state precisely what backward model is used, whether f is transformed, and how the noise is handled.

Moreover, Eq. (11) discretizes a reversed stochastic (difference) equation. It is known that (i) reversing an approximation introduces additional error, (ii) discretizing a reverse SDE to sufficient fidelity can be computationally very demanding even for moderate state dimensions, and (iii) there exist SDEs with bounded solutions whose discretizations are unbounded. Please discuss the discretization error, the time-step constraints, and the numerical stability of the backward pass, and comment on whether the asymptotic convergence claim survives these approximations.

Response:

Thank you for raising this important point! This comment is especially helpful for clarifying the conceptual interpretation of the EnKBS formulation. We agree that the backward EnKBS should not be interpreted as simply integrating the forward signal SDE backward in time. In general, the time reversal of a diffusion requires a density-dependent modification of the drift, as established by Anderson (1982). More directly related to the present setting, Anderson et al. (2021) construct an exact backward nonlinear smoothing diffusion whose drift depends on the filtering density and whose earlier-time marginals coincide with the smoothing distributions when initialized from the terminal filter law. In the additive-noise Gaussian case, their density-dependent term reduces to a covariance-weighted linear correction of the form

$$\Sigma P_f^{-1}(x - \bar{x}_f) = -\Sigma \nabla_x \log p_f(x),$$

where $\nabla_x \log p_f(x)$ is the Gaussian score. This has the same structural form as the smoothing correction appearing in the EnKBS. However, the present EnKBS is derived from the continuous-time limit of the ensemble RTS recursion rather than from this generic reverse-time diffusion construction. We do not replace f by a separately defined reverse drift. Instead, the backward evolution retains the model drift and adds the ensemble correction $\Sigma P_f^{-1}(x_s^{(i)} - x_f^{(i)})$. At the ensemble-member level, this differs from the Gaussian backward smoothing diffusion because the correction is centered at the paired filtering member $x_f^{(i)}$, rather than at the filtering mean. Likewise, the stochastic forcing in our construction is inherited from the forward ensemble trajectory: the same

Wiener realization associated with each filtering member is retained in the backward RTS-based recursion. Thus, Eq. (9) is not identical to the reverse-time diffusion of Anderson (1982) or the backward smoothing diffusion of Anderson et al. (2021). Nevertheless, in the linear–Gaussian mean-field limit as shown in Appendix B, the ensemble mean and covariance equations reduce to the classical continuous-time RTS smoother equations, consistent with the Gaussian specialization of Anderson et al. (2021). Importantly, the reuse of the forward noise realization is not just an implementation choice. Appendix A shows how it arises from the continuous-time limit of the ensemble RTS recursion, while Appendix B shows that this is used in recovering the classical RTS covariance equation.

Regarding the numerical discretization, we use Euler–Maruyama primarily for consistency and simplicity: the backward EnKBS is implemented using the same numerical scheme and the same time grid as the forward EnKBF, and this first-order backward discretization is consistent with the small-step expansion used in Appendix A. However, we agree that Eq. (11) introduces a finite time-step approximation and that Euler–Maruyama is not unconditionally stable for general nonlinear SDEs. In particular, for certain cases, explicit Euler approximations can exhibit numerical divergence or instability even when the underlying SDE remains well behaved. More robust schemes may alleviate such numerical-instability issues for broader classes of nonlinear SDEs, although at additional computational cost. In addition, there is also a distinct source of sensitivity associated with the backward propagation itself. Even a forward-stable system can amplify perturbations when propagated backward in time. In the EnKBS, the covariance-based smoothing correction couples each backward member to its stored filtering trajectory and acts to constrain this backward amplification, but this does not provide unconditional stability. The Lorenz-96 experiments illustrate this limitation: for small ensembles without sufficient localization, the backward pass can diverge even while the forward filter remains stable. Accordingly, Appendix A establishes the formal consistency between the discrete RTS recursion and the continuous-time EnKBS only in the vanishing time-step limit, while the exact RTS mean and covariance recovery in Appendix B applies to the continuous-time, linear–Gaussian EnKBS in the mean-field limit. At a finite time step and finite ensemble size, the practical implementation therefore contains additional discretization and finite-ensemble errors. In all numerical experiments, we use the same time step as in the forward EnKBF, sufficiently small to maintain numerical stability. A systematic numerical analysis establishing strong/weak convergence and step-size conditions for broad nonlinear SDE classes and alternative backward integrators is beyond the scope of the present paper, but will be an important direction for future work.

Major Comment 4: Offline versus online parameter estimation

Part of the current parameter-estimation literature relies on online, filter-based (state-augmentation) approaches, while the manuscript motivates offline parameter estimation; i.e., smoothing by the need for retrospective analysis once $y_{[0,T]}$ is available. The authors please explain why offline parameter estimation is preferable for the model-learning task.

Response:

Thank you for raising the distinction between online and retrospective parameter learning. We do not intend to suggest that offline parameter estimation is generally preferable to online, filter-based state-augmentation approaches. The two formulations serve different objectives: online methods are

appropriate when parameters need to be estimated or adapted sequentially as observations arrive. In contrast, in the model-learning problem considered here, the complete observation record $y_{[0,T]}$ is already available, and the learning algorithm is formulated retrospectively: at each iteration, the hidden state trajectory over the full window is sampled, and these trajectories are then used to update the model structure and parameters. Therefore, the EnKBS is particularly suitable for this setting because it incorporates the full observation record to reduce the uncertainty of the reconstructed hidden trajectory compared to the filter. We have revised Section 2 lines 134–136 and Section 3.3 lines 380–382 to clarify that this choice is specific to the present offline learning setup rather than a general superiority over online parameter-estimation methods.

Major Comment 5: Experimental choices and regularization

Several experimental choices would benefit from further clarification:

- (a) *Table 2 reports the best RMSE for $\delta^2 = 1.01$ and $r_0 = 4$, yet the time series in Figure 2 of x_1 from EnKBF/EnKBS are shown for $\delta^2 = 1.005$ and $r_0 = 3$. Please justify this choice or use the best-performing configuration.*
- (b) *Figure 3: the text states that inflation “does not noticeably change the results in this regime” and is therefore not applied, while elsewhere inflation is presented as needed to improve RMSE. Please reconcile these statements and specify in which regimes inflation is (not) required.*
- (c) *Section 3.1.2 mentions an ensemble size $m \leq 13$; this bound is unexplained. Please clarify.*

Response:

Thank you for carefully pointing out these experimental choices. We address the three points as follows:

(a) We have regenerated Figure 2 using the best-performing EnKBS configuration reported in Table 2, $(r_0, \delta^2) = (4, 1.01)$.

(b) Inflation is a regularization tool whose effect is problem-dependent rather than a universally required component of the EnKBS. In the high-dimensional Lorenz–96 experiment with a small ensemble, mild inflation can provide a modest reduction in RMSE, whereas localization has a substantially stronger effect on the stability of the backward smoother. In contrast, for the low-dimensional dyad regime considered in Section 3.2, applying mild inflation produces no noticeable change in the results (ACI metric and CIR patterns), and it is therefore not used. We have revised the Conclusion, lines 440–441 accordingly.

(c) The value 13 refers to the number of positive Lyapunov exponents of the deterministic Lorenz–96 system in the parameter regime considered here. We use $m = 10$, for which the empirical ensemble covariance has rank at most $m - 1 = 9 < 13$. Thus, the experiment deliberately considers an undersampled regime in which the ensemble covariance cannot fully resolve all unstable directions. We have clarified this point in Section 3.1.2 lines 283–285.

Major Comment 6: Observation model and generation of synthetic observations

No observation operator or observation model is specified for experimental models (e.g., Lorenz-96 in Section 3.1.1), so it is unclear how the ground truth of hidden state and observations were generated (which frequency, which noise distribution).

Response:

Thank you for pointing out this lack of clarity. Our observation setting differs from the conventional discrete-time observation model $y_k = Hx_k + \eta_k$. Instead, we partition the stochastic Lorenz-96 system into hidden and observed state variables: the odd-indexed components are treated as the hidden state, while the even-indexed components are treated as the observed state. Accordingly, the Lorenz-96 equations for the odd-indexed variables constitute the hidden-state dynamics in Eq. (1a), whereas the equations for the even-indexed variables constitute the continuous-time stochastic observation process in Eq. (1b). We have added a brief clarification at the beginning of Section 3, lines 253–256.

The full stochastic system is integrated to generate the reference trajectory, with mutually independent Wiener processes and noise variances $\sigma_h^2 = 5$ for the hidden components and $\sigma_o^2 = 0.1$ for the observed components. No additional discrete measurement-noise model is introduced. The observation interval equals the numerical integration time step, $\Delta t = 0.0005$, over $[0, T]$ with $T = 100$.

Major Comment 7: Novelty relative to existing ensemble smoothers

Novelty relative to existing ensemble smoothers: the first case study is presented as “highlighting the necessity of regularization and demonstrating the applicability of EnKBS” to a relatively high-dimensional chaotic system. Covariance localization and inflation are, however, standard approaches in EnKS and other extended ensemble smoothers, which already work with small ensemble sizes in high dimensions. The manuscript should state explicitly what the continuous-time formulation adds in the experiments (e.g., accuracy, stability for small ensemble size, cost, or theoretical convergence), ideally through a direct comparison with a discrete-time EnKS using the same regularization.

Response:

We appreciate this opportunity to better distinguish the continuous- and discrete-time formulations. The main contribution of the continuous-time formulation is conceptual and theoretical rather than a claimed improvement in accuracy, stability, or computational cost over a discrete-time EnKS. Many dynamical systems in scientific problems are naturally modeled in continuous time, often through stochastic differential equations. In such settings, the observations can also be represented as a continuous stochastic process rather than as a sequence of isolated measurements at discrete analysis times. This formulation is natural when the observed variables are themselves dynamically evolving components of the system and may interact with the hidden state, particularly for the causal inference and model discovery in this work. The continuous-time formulation also allows the EnKBS to serve as a natural backward counterpart to the EnKBF, so that the forward filtering and backward smoothing passes are formulated within the same continuous-time stochastic framework.

We also note that a direct numerical comparison with a discrete-time EnKS would not isolate the effect of the continuous-time formulation itself. The EnKBS in Eq. (9) is a continuous-time stochastic smoother, whereas the EnKS is defined in the discrete-time setting. Moreover, a discrete-time EnKS does not reduce to the EnKBS under a naive vanishing-step limit while keeping the observation-

noise covariance fixed; instead, the noise in the observation increments should be rescaled as $\sqrt{\Delta t}$, as shown in Appendix A. To perform such a numerical comparison, the EnKBS must first be discretized, as in Eq. (11). The resulting comparison would therefore be between two finite-step numerical schemes, and any differences in accuracy, stability, or computational cost would also depend on the chosen time-discretization scheme and step size, rather than solely on the distinction between continuous- and discrete-time smoothing. For this reason, we focus here on clarifying the role and interpretation of the continuous-time formulation rather than introducing an additional discrete-time benchmark.

Other Specific Comments

Specific Comment 1: Detection delay at extreme-event onset

Introduction: the claim that smoothing can “mitigates the detection delay of filters at the onset of extreme events” would benefit from a reference or the numerical evidence in the paper.

Response:

Thank you for pointing out the need for explicit support for this statement. The mechanism is illustrated directly by the nonlinear dyad experiment in Section 3.2. In the triggering direction $v \rightarrow u$, the hidden variable v controls the instantaneous growth rate of the observed variable u through $-d_u + cv$. The triggering condition is established once v becomes sufficiently large, but its effect becomes clearly visible in u only after the subsequent growth of u . A filter, which uses only observations up to the current time, can therefore exhibit a delay in identifying the hidden trigger before its effect is manifested in the observations. In contrast, the smoother uses the later growth of u to retrospectively improve the estimate of v during this preceding build-up phase.

This behavior is reflected in the numerical results in Section 3.2.3: the causal influence range becomes largest near the onset of v ’s growth, reflecting the large discrepancy between the filter and smoother before the corresponding extreme events of u become fully developed. We have revised the Introduction line 45 to point explicitly to this numerical evidence.

Specific Comment 2: Use of $dx^{(i)}$ in Eqs. (3) and (4)

Section 2: Equations (3) and (4) both use $dx^{(i)}$ for what appear to be different quantities (pseudo-time update with s versus physical-time increment with t). Please make this clear.

Response:

Thank you for pointing this out. This is discussed in the previous comment. We have removed this pseudo-time formulation and now introduce the stochastic EnKBF directly in physical time t .

Specific Comment 3: Noise notation in Lorenz-96

Section 3.1.1: two equations use the same symbol σ for different noise levels. Either distinguish them (σ_1, σ_2) or remove the redundant equation.

Response:

Thank you for pointing this out. We have revised Section 3.1.1 lines 274–276 to use distinct notation for the noise amplitudes of the hidden (odd-indexed) and observed (even-indexed) components, denoted by σ_h and σ_o , respectively.

Specific Comment 4: Definition of localization distance

Localization (page 8): the distance $d(i, j)$ depends on the state values x_i, x_j or on their indices i, j (e.g., cyclic index distance). If the latter, please correct the definition in the Localization subsection.

Response:

Thank you for pointing this out. Here, $d(i, j)$ denotes a distance between the indices i and j of the state components, rather than a distance depending on the state values x_i and x_j . We have reworded line 229 of page 8 in the "Localization and Inflation" subsection accordingly.

Specific Comment 5: Small-ensemble stability claims

Section 3.1.2: the statement that EnKBF remains stable down to $m = 3$ and that localized EnKBS remains stable down to $m = 5$ "(not shown)" is central to the small-ensemble claim of the abstract. Please show the EnKBF and EnKBS trajectories for both scenarios (with and without localization) so the reader can assess performance.

Response:

Thank you for raising this concern. To make the small-ensemble stability behavior directly assessable, we now provide the corresponding Lorenz–96 trajectories in Fig. 2. Further details on the parameter choices are given in the figure caption.

With localization and mild inflation, the EnKBF remains stable for $m = 3$, whereas the EnKBS becomes numerically unstable. We also found that varying the localization radius and inflation factor did not produce a stable EnKBS configuration at $m = 3$. Increasing the ensemble size to $m = 5$ stabilizes both the EnKBF and EnKBS. In the unlocalized and uninflated case, the EnKBF remains stable with $m = 17$, whereas the EnKBS is unstable. Increasing the ensemble size to $m = 31$ stabilizes both the EnKBF and EnKBS. The same qualitative behavior was reproduced across multiple independent random seeds and realizations. Figure 2 shows the realization corresponding to the same reference truth trajectory used in the manuscript.

We corrected the corresponding discussion in Section 3.1.2 lines 289–293.

Specific Comment 6: Operational definitions of ACI and CIR

Metrics (ACI and CIR): please give explicit formulas together with references, and state how the confidence/credible intervals used in these metrics are computed from the ensemble.

Response:




Figure 2: Lorenz-96 trajectories illustrating the effect of ensemble size and localization on the stability of the EnKBF and EnKBS. Apart from the ensemble size and the localization/inflation settings, all other model parameters and numerical settings are identical to those in the manuscript, with the same reference truth trajectory. Panels (a) and (b) use localization with $r_0 = 1$ and inflation $\delta^2 = 1.0001$, with ensemble sizes $m = 3$ and $m = 5$, respectively. For $m = 3$, the EnKBF remains stable while the EnKBS exhibits numerical instability, whereas for $m = 5$ both the filter and smoother remain stable. Panel (c) shows the unlocalized and uninflated case with $m = 17$: the EnKBF remains stable, whereas the EnKBS becomes numerically unstable and is therefore not displayed. Panel (d) shows the unlocalized and uninflated case with $m = 31$, for which both the EnKBF and EnKBS remain stable. The corresponding RMSEs are 1.003 (EnKBF) and 1.27×10^4 (EnKBS) for panel (a), 0.819 (EnKBF) and 0.714 (EnKBS) for panel (b), 1.069 (EnKBF) for panel (c), and 0.845 (EnKBF) and 0.800 (EnKBS) for panel (d). The choice $r_0 = 1$ corresponds to a relatively strong localization, consistent with the small ensemble size $m = 3$: with this localization radius, all more distant components outside of the central 3 components are truncated. We verified that setting $\delta^2 = 1$, i.e., using localization without inflation, leads to the same qualitative stability conclusions in panels (a) and (b). We also verified that for $m = 3$, varying r_0 and δ^2 does not yield a stable EnKBS configuration.

Thank you for this suggestion. Since the ACI and CIR are used here as diagnostics rather than the methodological focus of the paper, we retain the concise description in the manuscript and provide the explicit definitions here. Details and interpretations of the ACI and CIR are given in [Andreou et al. \(2026\)](#) and [Andreou and Chen \(2025\)](#).

Fix a time t in the time window $t \in [0, T]$. For a candidate causal variable $y(t)$ and observed trajectory x , the filtering distribution is

$$p_f(t) = p(y(t) \mid x(s \leq t)),$$

and the full smoothing distribution is

$$p_s(t) = p(y(t) \mid x(s \leq T)),$$

which incorporates the complete future observation. The instantaneous ACI metric is defined as the relative entropy between these two distributions,

$$\text{ACI} := \mathcal{P}(p_s(t), p_f(t)) = \int p_s(t) \log \frac{p_s(t)}{p_f(t)}.$$

To quantify the temporal extent of this causal influence, the full smoother is compared with lagged smoothing distributions

$$p(y(t) \mid x(s \leq T')), \quad t \leq T' \leq T,$$

which use future observations only up to T' . For a prescribed threshold $\epsilon > 0$, let T'_ϵ be the largest T' satisfying

$$\mathcal{P}(p(y(t) \mid x(s \leq T)), p(y(t) \mid x(s \leq T'))) > \epsilon.$$

The corresponding threshold-dependent causal inference range (CIR) is

$$\tau_\epsilon(t) := T'_\epsilon - t.$$

The objective CIR removes the dependence on a particular threshold by averaging over all admissible thresholds,

$$\tau(t) := \frac{1}{M(t)} \int_0^{M(t)} \tau_\epsilon(t) d\epsilon, \quad M(t) := \sup_{T' \in [t, T]} \mathcal{P}(p(y(t) \mid x(s \leq T)), p(y(t) \mid x(s \leq T'))).$$

We set $\tau(t) = 0$ when $M(t) = 0$. Dividing $M(t)$ guarantees the dimension of the objective CIR is ‘time’.

Regarding the confidence/credible intervals mentioned in the comment, such intervals are not defined for the ACI or CIR in the present framework and cannot be directly obtained from the ensemble spread. The ensemble is used to approximate the filtering, smoothing, and lagged-smoothing distributions through their empirical means and covariances, from which the ACI and CIR are subsequently computed. The ensemble members do not constitute independent realizations of the ACI or CIR; rather, they jointly approximate the conditional distribution associated with the same observed trajectory.

Additional Revision by the Authors

During the revision, we identified an issue in the original derivation in Appendix B involving an inappropriate application of the standard Itô covariance calculation to the fixed-interval smoothing process. We have therefore rewritten Appendix B to correct this. The corrected analysis shows that, in the linear–Gaussian mean-field limit, the EnKBS recovers the classical continuous-time RTS smoother mean and covariance equations when either the stochastic or deterministic EnKBF is used in the forward pass.

We have also revised Appendix A and the presentation of the backward EnKBS in Section 2.3. Several related notational and explanatory changes have been made throughout the manuscript for consistency.

We again thank the Editor and both Reviewers for their thoughtful and constructive comments.

Sincerely,

Zhang Jiang, Marios Andreou, Sebastian Reich, and Nan Chen

References

- B. D. O. Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982. [doi:10.1016/0304-4149\(82\)90051-5](https://doi.org/10.1016/0304-4149(82)90051-5).
- B. D. O. Anderson, A. N. Bishop, P. Del Moral, and C. Palmier. Backward nonlinear smoothing diffusions. *Theory of Probability and Its Applications*, 66(2):245–262, 2021. [doi:10.1137/S0040585X97T99037X](https://doi.org/10.1137/S0040585X97T99037X).
- M. Andreou, N. Chen, and E. Bollt. Assimilative causal inference. *Nature Communications*, 17:1854, 2026. [doi:10.1038/s41467-026-68568-0](https://doi.org/10.1038/s41467-026-68568-0).
- M. Andreou and N. Chen. Bridging Prediction and Attribution: Identifying Forward and Backward Causal Influence Ranges Using Assimilative Causal Inference. *arXiv preprint arXiv:2510.21889*, 2025. [doi:10.48550/arXiv.2510.21889](https://doi.org/10.48550/arXiv.2510.21889).
- D. Hodyss and M. Morzfeld. Using Diffusion Models to Do Data Assimilation. *Monthly Weather Review*, 154(2):165–182, 2026. [doi:10.1175/MWR-D-25-0125.1](https://doi.org/10.1175/MWR-D-25-0125.1).