



Ensemble Kalman–Guided Model Predictive Path Integral Control for Spatially Localized Suppression of Extremes in Chaotic Geophysical Flows

Haru Kuroki¹, Kazumune Hashimoto¹, Yuki Uehara¹, Yohei Sawada², Duc Le², and Masashi Minamide²

¹The University of Osaka

²The University of Tokyo

Correspondence: Kazumune Hashimoto (hashimoto@eei.eng.osaka-u.ac.jp)

Abstract. The possibility of influencing extreme weather phenomena has been discussed for decades; however, it remains far from operational practice, and there is still no established framework for designing small, spatially localized perturbations that can reliably steer chaotic geophysical flows. In this study, we propose a hybrid control method, termed ensemble-Kalman-guided model predictive path integral control (EKG-MPPI), which combines ensemble Kalman control (EnKC) with model predictive path integral (MPPI) control. Within a control simulation experiment framework, an ensemble Kalman filter is first used for state estimation, after which EnKC computes a candidate perturbation by treating the control objective as a pseudo-observation. An adaptive thresholding procedure then enforces spatial sparsity, so that the EnKC perturbation identifies candidate actuator locations and their nominal amplitudes. This information is embedded into the mean and covariance of Gaussian proposal distributions for MPPI, which subsequently refines the perturbation through sampling-based optimization with nonlinear rollouts, without linearizing the dynamics or computing gradients. Numerical experiments with the Lorenz-96 model and the surface quasi-geostrophic (SQG) model demonstrate that EKG-MPPI can suppress extremes in state variables and regional wind speed more effectively than EnKC alone, while using comparable or smaller control inputs. These results highlight EKG-MPPI as a promising building block for simulation-based assessment of localized intervention strategies in geophysical flows.

15 1 Introduction

The idea of deliberately influencing extreme weather phenomena, such as tropical cyclones (TCs), has been discussed for decades. However, it remains far from operational practice and there is still no established framework for designing small, spatially localized interventions with predictable effects. For example, Project STORMFURY attempted to weaken maximum wind speeds by artificially inducing convection around the TC eyewall, but its lack of effectiveness was reported by Willoughby et al. (1985). More recently, several studies have explored potential intervention mechanisms primarily in numerical modeling settings: based on simulations by Zhang et al. (2006), Cotton et al. (2007) discussed the possibility that mineral dust injection could suppress TC development, and Saharan dust has been shown to exert a strong influence on TCs. Furthermore, Latham et al. (2012) investigated TC weakening via sea-surface-temperature reduction, and Jacobson and Kempton (2014) demon-



strated in simulations that arrays of offshore wind turbines can suppress near-surface wind speeds. A comprehensive review of 25 these proposed approaches, together with feasibility and governance considerations, is provided by Miller et al. (2023). While these studies identify candidate actuators such as aerosols, mineral dust, and wind turbines, achieving concrete objectives (e.g., a substantial reduction in maximum wind speed) requires more than selecting an actuator type. It also calls for a mathematical framework that systematically optimizes *where*, *when*, and *how strongly* perturbations should be applied. This need is particularly acute because realistic anthropogenic influences are inherently local, weak, and intermittent. Therefore, it is crucial 30 to devise methods that can efficiently identify “small perturbations” that can nevertheless meaningfully steer the evolution of chaotic geophysical flows, at least in simulation-based assessment settings.

An early step in this direction was taken by Henderson, who applied four-dimensional variational data assimilation (4D-Var), widely used in numerical weather prediction, to derive optimal *initial* perturbations for mitigating damage from tropical cyclones. This approach assumes that perturbations are applied only at the initial time. In contrast, Sun et al. (2023) proposed a 35 framework in which small perturbations are applied continuously and adaptively. This concept, referred to as control simulation experiments (CSEs), has subsequently been employed and extended in later studies, including Kawasaki and Kotsuki (2024) and Ouyang et al. (2023), and has emerged as a practical approach for systematically evaluating the feasibility and potential side effects of candidate intervention strategies.

More recently, Sawada (2024a) introduced ensemble Kalman control (EnKC), which leverages the EnKF/EnKS framework 40 to compute control increments by assimilating the control objective as a pseudo-observation, and demonstrated its effectiveness in numerical experiments with the Lorenz-63 model. Building on this, Sawada (2024b) proposed a strategy to reduce the magnitude of EnKC control inputs and to enforce spatial sparsity, with demonstrations in the Lorenz-96 model. While these studies Sawada (2024a, b) indicate that EnKC can identify effective perturbations, its computation relies on an approximate linearization of the dynamics over the prediction horizon, motivating strategies that can more fully exploit nonlinear dynamics.

45 To address the linearization limitation of EnKC while retaining a derivative-free formulation, we combine EnKC with model predictive path integral (MPPI) control. Like EnKC, MPPI is derivative-free; however, it evaluates candidate perturbations through nonlinear forward rollouts rather than linearized dynamics. At the same time, MPPI can be sample-inefficient when informative prior knowledge about the control distribution is unavailable, particularly in high-dimensional control spaces Power and Berenson (2022). We therefore introduce EnKC-guided MPPI (EKG-MPPI), which uses the sparse perturbations obtained 50 from EnKC to construct an informative sampling distribution for MPPI. In this way, EnKC provides ensemble-based candidate perturbations that encode *where* control is likely to be effective and their nominal *magnitude*, while MPPI refines them via sampling-based optimization under the full nonlinear model without computing gradients. We demonstrate the effectiveness of the proposed method through numerical experiments with the Lorenz-96 model and the surface quasi-geostrophic (SQG) model.

55 This paper makes three contributions: (i) we propose EKG-MPPI, a hybrid framework that couples EnKC-based sparse actuation proposals with MPPI-based nonlinear refinement; (ii) we provide an explicit algorithmic formulation of EKG-MPPI, including a practical procedure for embedding EnKC-derived actuator locations and amplitudes into Gaussian sampling dis-



tributions for MPPI; and (iii) through numerical experiments with the Lorenz–96 and SQG models, we show that EKG-MPPI suppresses target extremes more strongly than EnKC (and uninformed MPPI) while using comparable or smaller control inputs.

60 2 Preliminary Knowledge

This section reviews the Ensemble Kalman filter/control and Model Predictive Path Integral control.

2.1 Ensemble Kalman Filter (EnKF)

Let us first review the EnKF. Consider a discrete-time state–space system

$$\mathbf{x}_t = M(\mathbf{x}_{t-1}) + \mathbf{q}_{t-1}, \quad (1)$$

$$65 \quad \mathbf{y}_t^o = H(\mathbf{x}_t) + \mathbf{r}_t, \quad (2)$$

where \mathbf{x}_t denotes the state vector, M the forecast model, \mathbf{q}_t the model error or the system noise, \mathbf{y}_t^o the observation vector, H the observation operator, and \mathbf{r}_t the observation error or the measurement noise. At an assimilation time t , the EnKF updates the model forecast by approximately minimizing the quadratic cost function

$$J(\mathbf{x}_t) = \frac{1}{2}(\mathbf{x}_t - \bar{\mathbf{x}}_t^b)^\top \mathbf{P}_b^{-1} (\mathbf{x}_t - \bar{\mathbf{x}}_t^b) + \frac{1}{2}(\mathbf{y}_t^o - H(\mathbf{x}_t))^\top \mathbf{R}^{-1} (\mathbf{y}_t^o - H(\mathbf{x}_t)), \quad (3)$$

70 where $\bar{\mathbf{x}}_t^b$ and \mathbf{P}_b are the ensemble mean and background error covariance, and \mathbf{R} is the observation error covariance matrix. The analysis ensemble $\{\mathbf{x}_t^{a(i)}\}_{i=1}^N$ is obtained by applying the EnKF update

$$\mathbf{x}_t^{a(i)} = \mathbf{x}_t^{f(i)} + \mathbf{K}(\mathbf{y}_t^{o(i)} - H(\mathbf{x}_t^{f(i)})), \quad (4)$$

$$\mathbf{K} = \mathbf{P}_b \mathbf{H}^\top (\mathbf{H} \mathbf{P}_b \mathbf{H}^\top + \mathbf{R})^{-1}, \quad (5)$$

where $\mathbf{x}_t^{f(i)}$ and $\mathbf{x}_t^{a(i)}$ denote the i th forecast and analysis ensemble members, respectively, and \mathbf{H} denotes the matrix representation of the (possibly linearized) observation operator H . In practice, the products $\mathbf{P}_b \mathbf{H}^\top$ and $\mathbf{H} \mathbf{P}_b \mathbf{H}^\top$ are computed using ensemble statistics, so that the full covariance matrix \mathbf{P}_b need not be formed explicitly. This ensemble representation enables the EnKF to handle high-dimensional systems efficiently. To mitigate sampling errors arising from the use of a finite ensemble, covariance localization is commonly employed. The localized Kalman gain is written as

$$\mathbf{K} = \rho \circ \mathbf{P}_b \mathbf{H}^\top (\rho \circ (\mathbf{H} \mathbf{P}_b \mathbf{H}^\top) + \mathbf{R})^{-1}, \quad (6)$$

80 where \circ denotes the Schur (element-wise) product and

$$\rho = \exp[-d(i, j)/L] \quad (7)$$

is a smooth correlation function defined in terms of the distance $d(i, j)$ between grid points and a localization length scale L . This localization suppresses spurious long-range correlations, thereby improving the stability and accuracy of the EnKF updates.



85 Closely related to the EnKF is the ensemble Kalman smoother (EnKS), which extends the filter in time by estimating the state over an assimilation window using both past and future observations. Operationally, the EnKS can be implemented by propagating an ensemble forward with EnKF updates at each observation time and then applying a backward smoothing step that uses ensemble-based cross-covariances between states at different times and the observations. Readers interested in further details of the EnKF/EnKS family are referred to Houtekamer and Zhang (2016) for a comprehensive review.

90 **2.2 Ensemble Kalman Control (EnKC)**

EnKC builds directly on the EnKF and EnKS framework to formulate and solve an optimal control problem for various state-space (dynamical) systems. After the EnKF analysis at time t , let $\{\mathbf{x}_t^{a(i)}\}_{i=1}^N$ denote the analysis ensemble, with mean $\bar{\mathbf{x}}_t^a$ and error covariance \mathbf{P}_a . Over a prediction horizon T_{EnKC} , EnKC considers the quadratic cost function

$$J_c(\mathbf{x}_t) = \frac{1}{2} (\mathbf{x}_t - \bar{\mathbf{x}}_t^a)^\top \mathbf{P}_a^{-1} (\mathbf{x}_t - \bar{\mathbf{x}}_t^a) + \frac{1}{2} (\mathbf{r}_{t+T_{\text{EnKC}}} - H_c(\mathbf{x}_{t+T_{\text{EnKC}}}))^\top \mathbf{R}_c^{-1} (\mathbf{r}_{t+T_{\text{EnKC}}} - H_c(\mathbf{x}_{t+T_{\text{EnKC}}})), \quad (8)$$

95 subject to the model dynamics

$$\mathbf{x}_{k+1} = M(\mathbf{x}_k), \quad k = t, \dots, t + T_{\text{EnKC}} - 1. \quad (9)$$

Here, $\mathbf{r}_{t+T_{\text{EnKC}}}$ is a prescribed reference (target) vector at time $t + T_{\text{EnKC}}$, H_c maps the state variables to the control criteria (e.g., a regional average or maximum of a physical quantity), and \mathbf{R}_c is a user-defined weight (pseudo-observation error covariance). The first term penalizes the size of the perturbation to the initial analysis state, while the second term penalizes the mismatch between the predicted future state and the control objective. Assuming that the dynamics over the prediction horizon are approximately linear, the minimizer of J_c can be obtained by applying an EnKS. In this formulation, the reference vector $\mathbf{r}_{t+T_{\text{EnKC}}}$ is assimilated as a “pseudo-observation” with error covariance \mathbf{R}_c . The resulting EnKS analysis at time t ,

$$\mathbf{x}_t^c = \bar{\mathbf{x}}_t^a + K(\mathbf{r}_{t+T_{\text{EnKC}}} - H_c(\mathbf{x}_{t+T_{\text{EnKC}}^a}^a)), \quad (10)$$

105 yields the optimal perturbation $\mathbf{x}_t^c - \bar{\mathbf{x}}_t^a$ to be applied to the system. The Kalman gain K is computed from the cross-covariance between the analysis ensemble at time t and the ensemble prediction at $t + T_{\text{EnKC}}$, and from the covariance of the projected prediction, in direct analogy with standard EnKF/EnKS formulations.

The EnKC algorithm can be summarized as follows:

1. Apply EnKF with real observations to obtain the analysis ensemble at time t .
2. From the analysis ensemble, perform ensemble forecasting up to $t + T_{\text{EnKC}}$ and project the predicted states onto the control criteria via H_c .
- 110 3. Run EnKS, assimilating $\mathbf{r}_{t+T_{\text{EnKC}}}$ as a pseudo-observation with error covariance \mathbf{R}_c , and obtain the perturbation $\mathbf{x}_t^c - \bar{\mathbf{x}}_t^a$.
4. Add this perturbation to the real system and to all analysis ensemble members, thereby updating the controlled “nature” and its ensemble representation.



5. Propagate the updated ensemble to the next data assimilation time. The resulting forecast serves as the new background
 115 for the next EnKF cycle, thus completing the loop and returning to Step 1.

In this way, EnKC integrates ensemble data assimilation and model predictive control by treating the control objective as a pseudo-observation within an EnKS framework, and by interpreting the resulting analysis increment as the optimal small perturbation to steer the system toward the desired future state. The control perturbation estimated by EnKC (Sawada (2024b)),
 $\mathbf{x}_t^c - \bar{\mathbf{x}}_t^a$, has the same dimension as the full model state. In practical applications such as weather modification, however, it is
 120 unrealistic to apply perturbations to all state variables at every control step. It is therefore desirable to enforce *spatial sparsity* so that only a limited number of grid points (or regions) are actively perturbed, which naturally aligns with the interpretation of actuators placed at specific locations. Ideally, such sparsity can be promoted by augmenting the EnKC cost function with an
 ℓ_0 -norm penalty:

$$J_c(\mathbf{x}_t) = \frac{1}{2}(\mathbf{x}_t - \bar{\mathbf{x}}_t^a)^\top \mathbf{P}_a^{-1}(\mathbf{x}_t - \bar{\mathbf{x}}_t^a) + \frac{1}{2}(\mathbf{r}_{t+T_{\text{EnKC}}} - H_c(\mathbf{x}_{t+T_{\text{EnKC}}}))^\top \mathbf{R}_c^{-1}(\mathbf{r}_{t+T_{\text{EnKC}}} - H_c(\mathbf{x}_{t+T_{\text{EnKC}}})) + \lambda_s \|\mathbf{x}_t - \bar{\mathbf{x}}_t^a\|_0, \quad (11)$$

125 where $\|\cdot\|_0$ counts the number of nonzero elements (i.e., ℓ_0 -norm) and λ_s controls the strength of the sparsity constraint. However, direct minimization of this objective is computationally demanding for high-dimensional systems, so an empirical yet effective strategy is adopted.

Following the idea of Schneider et al. (2022), we impose sparsity on the EnKC-estimated perturbation by applying a thresholding operator to its components. Specifically, after computing the standard EnKC update, we apply a thresholding operator
 130 $T(\theta)$ to each entry of $\mathbf{x}_t^c - \bar{\mathbf{x}}_t^a$ so that small-magnitude components are discarded as noise and only sufficiently large perturbations are retained:

$$T(\theta) = \begin{cases} 0, & \text{if } |\theta| < \sqrt{2\lambda_s}, \\ \theta, & \text{otherwise.} \end{cases} \quad (12)$$

While Schneider et al. (2022) treated $\sqrt{2\lambda_s}$ as a fixed hyperparameter, the present study adopts an *adaptive* formulation in which the threshold depends on the amplitude of the EnKC-derived perturbation at each control step. More precisely, the
 135 threshold parameter is updated as

$$\sqrt{2\lambda_s} = \Lambda \max |\mathbf{x}_t^c - \bar{\mathbf{x}}_t^a|. \quad (13)$$

where Λ is a user-specified coefficient that controls the enforced sparsity. This adaptive thresholding allows the sparsity level to automatically adjust to the magnitude of the EnKC-derived perturbations, thereby concentrating actuator effort on only the most influential grid points. In particular, choosing $\Lambda = 1$ retains only the grid point with the largest perturbation, which is
 140 consistent with the actuator-placement interpretation in which localized increments indicate candidate actuator locations.



2.3 Model Predictive Path Integral Control (MPPI)

We next summarize the MPPI control framework. MPPI is a sample-based model predictive control (MPC) method grounded in probabilistic inference (Williams et al. (2018)). MPPI considers the following nonlinear dynamical system:

$$\mathbf{x}_{t+1} = M(\mathbf{x}_t, \mathbf{u}_t). \quad (14)$$

145 At each time step, MPPI optimizes an *open-loop* control sequence over a finite horizon of length T . Let $\mathbf{U} = (\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{T-1})$ denote the *mean* control sequence. MPPI introduces stochastic exploration by sampling control sequences $\mathbf{V} = (\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_{T-1})$ from a Gaussian distribution centered at \mathbf{U} :

$$q(\mathbf{V} \mid \mathbf{U}, \Sigma) = \prod_{\tau=0}^{T-1} \mathcal{N}(\mathbf{v}_\tau; \mathbf{u}_\tau, \Sigma), \quad (15)$$

150 where $\Sigma \in \mathbb{R}^{n_u \times n_u}$ is a user-specified positive definite covariance matrix. Equivalently, one may write $\mathbf{v}_\tau = \mathbf{u}_\tau + \boldsymbol{\epsilon}_\tau$ with $\boldsymbol{\epsilon}_\tau \sim \mathcal{N}(\mathbf{0}, \Sigma)$. Given a sampled control sequence \mathbf{V} , the system is rolled out according to $\mathbf{x}_{\tau+1} = M(\mathbf{x}_\tau, \mathbf{v}_\tau)$ (starting from the current state), and we define the total trajectory cost functional as

$$S(\mathbf{V}) = \Phi(\mathbf{x}_T) + \sum_{\tau=0}^{T-1} L(\mathbf{x}_\tau, \mathbf{v}_\tau). \quad (16)$$

Under this setup, the stochastic optimal control problem at each time step is

$$\mathbf{U}^* = \arg \min_{\mathbf{U} \in \mathcal{U}} \mathbb{E}_{\mathbf{V} \sim q(\mathbf{V} \mid \mathbf{U}, \Sigma)} [S(\mathbf{V})]. \quad (17)$$

155 MPPI can be derived via a variational free-energy bound. Let $p(\mathbf{V})$ denote a *base* (prior) distribution over control sequences (typically chosen as a Gaussian), and define the free energy

$$F = -\lambda \log \mathbb{E}_{\mathbf{V} \sim p} \left[\exp \left(-\frac{S(\mathbf{V})}{\lambda} \right) \right]. \quad (18)$$

Then, for any distribution $r(\mathbf{V})$, the following inequality holds:

$$F \leq \mathbb{E}_{\mathbf{V} \sim r} [S(\mathbf{V})] + \lambda D_{\text{KL}}(r \parallel p), \quad (19)$$

160 where D_{KL} is the Kullback–Leibler divergence. The minimizer of the right-hand side of (19) is the optimal distribution

$$q^*(\mathbf{V}) = \frac{1}{\eta} \exp \left(-\frac{S(\mathbf{V})}{\lambda} \right) p(\mathbf{V}), \quad (20)$$

and substituting $r(\mathbf{V}) = q^*(\mathbf{V})$ into (19) yields equality (see Williams et al. (2018) for details). In practice, MPPI restricts $r(\mathbf{V})$ to a Gaussian family with fixed covariance Σ and optimizes only the mean \mathbf{U} . This corresponds to the KL projection

$$\hat{\mathbf{U}} = \arg \min_{\mathbf{U} \in \mathcal{U}} D_{\text{KL}}(q^* \parallel q(\mathbf{V} \mid \mathbf{U}, \Sigma)), \quad (21)$$



165 whose minimizer satisfies, for fixed Σ ,

$$\hat{\mathbf{U}} = \mathbb{E}_{\mathbf{V} \sim q^*} [\mathbf{V}]. \quad (22)$$

Since direct sampling from q^* is intractable, we estimate (22) by Monte Carlo sampling with importance weights. When samples $\{\mathbf{V}_k\}_{k=1}^K$ are drawn from the base distribution $p(\mathbf{V})$, (22) can be written as a normalized weighted average:

$$\hat{\mathbf{U}} \approx \sum_{k=1}^K \bar{\omega}(\mathbf{V}_k) \mathbf{V}_k, \quad (23)$$

170 with log-weights and normalized weights

$$\omega(\mathbf{V}_k) = -\frac{1}{\lambda} S(\mathbf{V}_k), \quad (24)$$

$$\bar{\omega}(\mathbf{V}_k) = \frac{\exp(\omega(\mathbf{V}_k))}{\sum_{j=1}^K \exp(\omega(\mathbf{V}_j))}. \quad (25)$$

(If samples are drawn from a proposal distribution $g(\mathbf{V}) \neq p(\mathbf{V})$, the weights are modified by the standard importance-sampling factor $p(\mathbf{V}_k)/g(\mathbf{V}_k)$.)

175 Typically, only the first element of the optimized control sequence $\hat{\mathbf{U}}$ is applied as the control input at each time step, after which the horizon is shifted forward and the procedure is repeated. As can be seen from (23)–(25), the weights for all samples can be computed in parallel, allowing efficient implementation on GPUs. Moreover, MPPI does not require gradient information of the dynamics or cost function and can be implemented using only forward simulations of the model.

3 EKG-MPPI: Proposed hybrid control method

180 3.1 Overall framework

EnKC provides a principled framework for determining small yet effective perturbations at each control step. However, EnKC computes the control increment by (approximately) linearizing the system dynamics over a finite prediction horizon, and its performance may therefore deteriorate when the dynamics are strongly nonlinear. To overcome this limitation, we propose a hybrid control scheme, termed *EnKC-guided MPPI* (EKG-MPPI), which combines EnKC with MPPI. In contrast to EnKC, 185 MPPI evaluates candidate perturbations via nonlinear forward rollouts and thus can handle nonlinear dynamics without linearization. Nevertheless, as pointed out by Power and Berenson (2022), MPPI may suffer from poor sampling efficiency when applied without informative prior knowledge of the control distribution. This issue is particularly severe in high-dimensional geophysical systems, where the control space is large and naive sampling may fail to discover effective perturbations under a realistic computational budget.

190 EKG-MPPI addresses this issue by using the EnKC perturbation as *prior information* for MPPI. Intuitively, EnKC provides a sparse and physically informed guess of (i) *where* the system is most sensitive and (ii) *how strongly* it should be perturbed. We embed this information into the mean and variance (or covariance) of Gaussian sampling distributions, from which MPPI

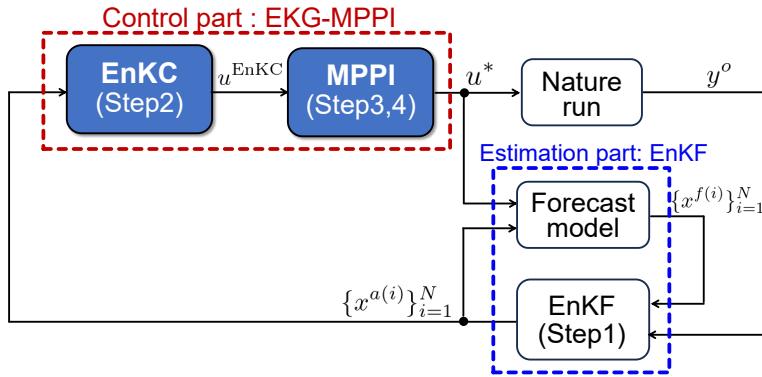


Figure 1. Block diagram of the proposed EKG-MPPI.

draws candidate actuator locations and magnitudes. MPPI then refines the perturbation through sampling-based optimization based on nonlinear rollouts, without any linearization or gradient computation.

195 The proposed EKG-MPPI scheme is implemented within the CSE framework, and the overall workflow is summarized in Fig. 1. At each data-assimilation time, we first apply the EnKF to assimilate observations and estimate the current state (Step 1). We then run EnKC using the analysis ensemble to obtain a sparse perturbation (Step 2). This perturbation is interpreted as a candidate actuator configuration and mapped to the parameters of Gaussian sampling distributions for actuator location and control magnitude. MPPI samples multiple candidate control inputs from these distributions, evaluates their performance over 200 a prediction horizon, and computes a weighted average to obtain the final control input (Steps 3–4). In the CSE framework, the model trajectory obtained by integrating the forecast model without any control perturbation is referred to as the *nature run*, which serves as a proxy for the true atmosphere. The computed control is applied both to the nature run and to the analysis ensemble, after which the forecast–assimilation–control loop continues to the next time step.

3.2 Algorithmic steps (Step 1–5)

205 In this subsection we describe the EKG-MPPI procedure at a single control time t . We denote the ensemble size by N and the current analysis ensemble by $\{x_t^{a(i)}\}_{i=1}^N$.

Step 1: EnKF state estimation

As in standard CSE studies, accurate estimation of the system state is a prerequisite for effective control. At each assimilation time t , we apply the EnKF to assimilate the available observations y_t^o and update the forecast ensemble. The EnKF update is



210 given in Sect. 2.1 and yields the analysis ensemble $\{x_t^{a(i)}\}_{i=1}^N$. We then compute the analysis ensemble mean

$$\bar{x}_t^a = \frac{1}{N} \sum_{i=1}^N x_t^{a(i)}, \quad (26)$$

which serves as the current state estimate used to initialize the subsequent control computation.

Step 2: EnKC-based control perturbation

215 In the second step, we apply EnKC (Sect. 2.2) to compute a control perturbation that will serve as prior information for MPPI.
 EnKC minimizes the quadratic cost function in (8) via an EnKS, interpreting the control objective as a pseudo-observation. The resulting control perturbation $x_t^c - \bar{x}_t^a$ is then sparsified by the adaptive thresholding procedure described in Sect. 2.2. Among the information contained in the sparsified control vector, the key components for localized intervention are (i) the indices of its nonzero entries, which indicate candidate actuator locations, and (ii) the corresponding magnitudes, which represent nominal control amplitudes at those locations. To simplify the presentation, we set $\Lambda = 1$ in (13), so that only the grid point with the 220 largest perturbation remains nonzero. The resulting sparse control vector is denoted by u_t^{EnKC} , and encodes both the actuator location and its nominal control magnitude to be exploited by MPPI.

Step 3: EnKC-informed sampling of actuator location and magnitude

225 The third step embeds the information contained in u_t^{EnKC} into the sampling distributions from which MPPI draws candidate control inputs. In this study, we parameterize each candidate actuator configuration by its location and magnitude. Let $\mu_{\text{loc},t}$ and $\Sigma_{\text{loc},t}$ denote the mean and variance (scalar in the single-actuator case) that determine the sampling distribution of actuator locations, and let $\mu_{\text{mag},t}$ and $\Sigma_{\text{mag},t}$ denote the corresponding quantities for the control magnitude. We define

$$\begin{aligned} \mu_{\text{loc},t} &= f_{\text{loc}}(u_t^{\text{EnKC}}), & \Sigma_{\text{loc},t} &= g_{\text{loc}}(u_t^{\text{EnKC}}), \\ \mu_{\text{mag},t} &= f_{\text{mag}}(u_t^{\text{EnKC}}), & \Sigma_{\text{mag},t} &= g_{\text{mag}}(u_t^{\text{EnKC}}), \end{aligned} \quad (27)$$

230 where f_{loc} , g_{loc} , f_{mag} , and g_{mag} are embedding functions that map the EnKC output to the parameters of the MPPI sampling distributions. Using the parameters in (27), we draw K_{MPPI} samples of actuator locations and magnitudes,

$$l_{t,i} \sim \mathcal{G}(l \mid \mu_{\text{loc},t}, \Sigma_{\text{loc},t}), \quad m_{t,i} \sim \mathcal{G}(m \mid \mu_{\text{mag},t}, \Sigma_{\text{mag},t}), \quad i = 1, 2, \dots, K_{\text{MPPI}}, \quad (28)$$

where $\mathcal{G}(\cdot \mid \mu, \Sigma)$ denotes a Gaussian distribution. Because actuator locations correspond to discrete grid indices, we round $l_{t,i}$ to the nearest integer and clip it to the valid index range; with a slight abuse of notation, we denote the resulting index again by $l_{t,i}$. Each pair $(l_{t,i}, m_{t,i})$ specifies one candidate actuator configuration for MPPI.

235 Step 4: MPPI rollout and importance weighting

In the fourth step, we evaluate the sampled actuator configurations via MPPI rollouts and compute their importance weights. Let $x_t = \bar{x}_t^a$ denote the current state estimate (i.e., the state from which MPPI rolls out). Let n_u denote the dimension of the



control input, and let $e_\ell \in \mathbb{R}^{n_u}$ be the ℓ -th standard basis vector. For each sample $(l_{t,i}, m_{t,i})$, we construct a sparse control vector

240 $u_{t,i} = m_{t,i} e_{l_{t,i}}.$ (29)

Starting from $x_{t,0}^{(i)} = x_t$, we propagate the system over the MPPI prediction horizon T_{MPPI} as

$$x_{t,1}^{(i)} = M(x_{t,0}^{(i)}, u_{t,i}), \quad (30)$$

$$x_{t,\tau+1}^{(i)} = M(x_{t,\tau}^{(i)}, \mathbf{0}), \quad \tau = 1, 2, \dots, T_{\text{MPPI}} - 1, \quad (31)$$

where $\mathbf{0}$ denotes the zero control input. That is, the perturbation is applied only at the first rollout step, and no further input 245 is provided during the remaining rollout. This reflects the single-step actuation setting used in our CSEs and also reduces the effective dimension of the control space, thereby improving sampling efficiency.

Let $S(u_{t,i}, x_t, M)$ be a state-dependent cost functional that quantifies the performance of the i -th control sample. In our implementation, the EnKC-informed Gaussian distribution in Step 3 is used as the base distribution of MPPI, and therefore no importance-sampling correction term is required. The (unnormalized) log-weight is computed as

250 $\omega(u_{t,i}) = -\frac{1}{\lambda} S(u_{t,i}, x_t, M),$ (32)

where λ is the temperature parameter in MPPI. The normalized weight is then computed as

$$\bar{w}_{t,i} = \frac{\exp(\omega(u_{t,i}))}{\sum_{k=1}^{K_{\text{MPPI}}} \exp(\omega(u_{t,k}))}. \quad (33)$$

Finally, we first compute the weighted-average control vector

$$\tilde{u}_t = \sum_{i=1}^{K_{\text{MPPI}}} \bar{w}_{t,i} u_{t,i}, \quad (34)$$

255 and then project it to a single-actuator (1-sparse) control by keeping only the largest-magnitude component:

$$l_t^* = \arg \max_{\ell} |\tilde{u}_t(\ell)|, \quad (35)$$

$$u_t^* = \tilde{u}_t(l_t^*) e_{l_t^*}. \quad (36)$$

Step 5: Feedback

260 In the final step, the optimal control u_t^* computed in (36) is applied both to the nature run and to each member of the analysis ensemble, i.e.

$$x_t^{a(i)} \leftarrow x_t^{a(i)} + u_t^*, \quad i = 1, 2, \dots, N. \quad (37)$$

The updated ensemble is then advanced by the forecast model until the next assimilation time, at which point the procedure returns to Step 1.



Algorithm 1 EKG-MPPI algorithm

Require: initial analysis ensemble $\{\mathbf{x}_0^{a(i)}\}_{i=1}^N$; localization scale L ; observation-error covariance matrix \mathbf{R} ; forecast model M ; observation operator H ; control-weight matrix \mathbf{R}_c ; EnKC prediction horizon T_{EnKC} ; MPPI prediction horizon T_{MPPI} ; sparsity coefficient Λ ; state-cost functional S ; embedding functions $f_{\text{loc}}, g_{\text{loc}}, f_{\text{mag}}, g_{\text{mag}}$; number of MPPI samples K_{MPPI} ; temperature parameter λ ; total number of assimilation steps T_{sim} .

- 1: **for** $t = 1, 2, \dots, T_{\text{sim}}$ **do**
- 2: Obtain observation \mathbf{y}_t^o .
- 3: $\{\mathbf{x}_t^{a(i)}\}_{i=1}^N \leftarrow \text{EnKF}(\{\mathbf{x}_{t-1}^{a(i)}\}_{i=1}^N, \mathbf{y}_t^o, L, \mathbf{R}, M, H)$.
- 4: Compute the analysis ensemble mean $\bar{\mathbf{x}}_t^a \leftarrow \frac{1}{N} \sum_{i=1}^N \mathbf{x}_t^{a(i)}$.
- 5: $u_t^{\text{EnKC}} \leftarrow \text{EnKC}(\{\mathbf{x}_t^{a(i)}\}_{i=1}^N, T_{\text{EnKC}}, \Lambda, \mathbf{R}_c, M, H_c)$.
- 6: Compute $\mu_{\text{loc},t}, \Sigma_{\text{loc},t}, \mu_{\text{mag},t}, \Sigma_{\text{mag},t}$ according to (27).
- 7: **for** $i = 1, 2, \dots, K_{\text{MPPI}}$ **do**
- 8: Sample $l_{t,i} \sim \mathcal{G}(l \mid \mu_{\text{loc},t}, \Sigma_{\text{loc},t})$, $m_{t,i} \sim \mathcal{G}(m \mid \mu_{\text{mag},t}, \Sigma_{\text{mag},t})$.
- 9: Construct a sparse control vector $u_{t,i} \leftarrow m_{t,i} e_{l_{t,i}}$, where e_ℓ denotes the ℓ -th standard basis vector.
- 10: Evaluate the cost $S(u_{t,i}, \bar{\mathbf{x}}_t^a, M)$ and compute the (unnormalized) log-weight $\omega(u_{t,i})$ using (32).
- 11: **end for**
- 12: Normalize the weights $\bar{w}_{t,i}$ using (33).
- 13: Compute the weighted-average control vector $\tilde{u}_t \leftarrow \sum_{i=1}^{K_{\text{MPPI}}} \bar{w}_{t,i} u_{t,i}$.
- 14: Select the actuator location $l_t^* \leftarrow \arg \max_\ell |\tilde{u}_t(\ell)|$.
- 15: Set the single-actuator control $u_t^* \leftarrow \tilde{u}_t(l_t^*) e_{l_t^*}$.
- 16: Apply u_t^* to the nature run and update each analysis ensemble member via $\mathbf{x}_t^{a(i)} \leftarrow \mathbf{x}_t^{a(i)} + u_t^*, i = 1, \dots, N$.
- 17: Advance the forecast model M to the next assimilation time.
- 18: **end for**

265 **3.3 Pseudocode of EKG-MPPI**

The EKG-MPPI procedure described above can be summarized by Algorithm 1, where **EnKF** (line 3) denotes the ensemble Kalman filter used for state estimation and **EnKC** (line 5) denotes the ensemble Kalman control scheme that computes sparse optimal perturbations as described in Sect. 2.2.

4 Numerical Experiments

270 The effectiveness of EKG-MPPI is demonstrated through numerical experiments using the Lorenz–96 model and the SQG model.



4.1 Lorenz-96 model

Following the methodology of Sun et al. (2023); Sawada (2024b), we conduct a control simulation experiment (CSE) aimed at mitigating extreme values in the Lorenz-96 system (Lorenz (1995)). The Lorenz-96 system is governed by

$$275 \quad \frac{dX_k}{dt} = (X_{k+1} - X_{k-2})X_{k-1} - X_k + F, \quad k \in \{1, 2, \dots, K\}, \quad (38)$$

where $K = 40$ in this study, and the external forcing parameter is set to $F = 8.0$. Cyclic boundary conditions are imposed such that the indices are taken modulo K , i.e. $X_{k+K} = X_k$. Equation (38) is numerically integrated using the fourth-order Runge–Kutta method with a time step of $\Delta t = 0.05$.

Observations are available only at grid points with even indices (20 state variables) and are generated by adding Gaussian 280 white noise with mean 0 and variance 1.0 to the nature run. The observation interval is set to 0.05, i.e., observations are taken at every model time step. The ensemble size is $N = 40$, and the observation-error covariance matrix is $\mathbf{R} = \mathbf{I}$ with unit variance. The EnKF is applied at each observation time with a localization length of $L = 2.0$ (see (7)). The control objective is to suppress extremely large positive anomalies in the system. At each EnKF analysis time t , an additional forecast is integrated 285 over a prediction horizon of $T_c = 0.2$, corresponding to four model time steps. EnKC is activated only when the ensemble-mean forecast at $t + T_c$ indicates $X_k > 12$ at one or more grid points. Following Sun et al. (2023); Sawada (2024b), we interpret $T_c = 0.2$ time units (four model steps) as a short lead time and use $X_k = 12$ as a high-percentile threshold for extremes in the Lorenz-96 system. Since the control goal is to prevent further growth beyond this threshold, the target pseudo-observation is set to $r_{t+T_c} = 12$. In the EnKC procedure, this pseudo-observation is assimilated using the EnKS, but only at grid points where $X_k > 12$. The diagonal entries of the control-weight matrix \mathbf{R}_c are set to 10^{-4} . The sparse control perturbation computed 290 by EnKC, denoted by $\mathbf{u}_t^{\text{EnKC}}$, is then used to construct the MPPI prior in EKG-MPPI. The embedding functions in (27) are specified as

$$f_{\text{loc}}(u_t^{\text{EnKC}}) = \text{argnz}(u_t^{\text{EnKC}}), \quad g_{\text{loc}}(u_t^{\text{EnKC}}) = 5.0, \\ f_{\text{mag}}(u_t^{\text{EnKC}}) = u_t^{\text{EnKC}}(f_{\text{loc}}(u_t^{\text{EnKC}})), \quad g_{\text{mag}}(u_t^{\text{EnKC}}) = \frac{\|u_t^{\text{EnKC}}\|_1}{2}, \quad (39)$$

where $\text{argnz}(\cdot)$ returns the index of the nonzero entry of its argument. For the initial solution of MPPI, we adopt the perturbations obtained by EnKC, as represented by f_{loc} and f_{mag} . The variance of the Gaussian distribution used for sampling is 295 adaptively adjusted according to g_{mag} . Furthermore, control samples whose magnitude exceeds the maximum absolute component of u_t^{EnKC} are rejected, ensuring that the applied control remains sufficiently small. For MPPI, the prediction horizon is set to $T_{\text{MPPI}} = 5$, and the running state-cost function is defined as

$$S(x_t) = \|\max(x_t - 12, 0)\|_1, \quad (40)$$

which penalizes the state only when it exceeds the threshold value of 12. We perform numerical simulations over 160,600 model 300 time steps. During the first 14,600 time steps, neither EnKC nor EKG-MPPI is applied; only EnKF is used to synchronize the estimated state with the uncontrolled nature run (spin-up period). Control is activated for the remaining 146,000 time steps,



during which EnKF, EnKC, and EKG-MPPI are all executed. All diagnostics reported below are computed over this controlled period.

4.1.1 Simulation Results

305 To demonstrate the effectiveness of the proposed method, we compare four cases—no control (uncontrolled), EnKC, vanilla MPPI, and EKG-MPPI—in terms of the state distribution during the simulation and the control effort required to suppress extremes. For EKG-MPPI, we set the temperature parameter to $\lambda = 0.1$ and the number of control-input samples to $K_{\text{MPPI}} = 40$. For vanilla MPPI, we set $\lambda = 0.7$ and $K_{\text{MPPI}} = 40$; the perturbation magnitude is sampled from a Gaussian distribution $\mathcal{N}(0, 0.5)$, and the perturbation location is sampled from $\mathcal{N}(0, 20.0)$.

310 The simulation results are summarized in Figure 2 and Table 1. Figure 2a shows that the percentile curves over the full range remain close to the uncontrolled case for all controllers, indicating that the interventions primarily affect the upper tail while leaving the bulk distribution nearly unchanged in this experiment. The effect of control is concentrated in the far upper tail (Figure 2b), where all controllers reduce the most extreme values relative to the uncontrolled run.

315 Table 1 quantifies extreme-event suppression and control effort. EKG-MPPI reduces the number of exceedances at $X \geq 12.5$ and $X \geq 13.0$ compared with both EnKC ($3293 \rightarrow 3215$ and $1308 \rightarrow 1176$) and vanilla MPPI ($3528 \rightarrow 3215$ and $1284 \rightarrow 1176$). At the same time, EKG-MPPI achieves these reductions with smaller control inputs: the maximum input decreases from 1.311 (EnKC) and 1.780 (MPPI) to 1.037, and the mean input decreases from 1.638×10^{-4} (EnKC) and 1.547×10^{-4} (MPPI) to 1.187×10^{-4} . Here the mean input is computed as a time- and grid-averaged absolute input (including time steps with no control where $u_{t,k} = 0$). We attribute the improvement over EnKC to the nonlinear evaluation in the MPPI refinement stage, 320 and the improvement over vanilla MPPI to the EnKC-informed proposal distribution, which concentrates samples on effective actuator locations and magnitudes under a fixed sampling budget.

Table 1. Comparison of EKG-MPPI, EnKC, and MPPI in the Lorenz-96 control simulation experiment. Columns 2–3 quantify the control effort: Max Input = $\max_{t,k} |u_{t,k}|$ and Mean Input = $\frac{1}{TK} \sum_{t=1}^T \sum_{k=1}^K |u_{t,k}|$ (time- and grid-averaged absolute input, including time steps with no control where $u_{t,k} = 0$). Columns 4–5 report extreme-event counts, defined as $\#\{(t, k) : X_{t,k} \geq \theta\}$ for thresholds $\theta \in \{12.5, 13.0\}$, where $\#\{\cdot\}$ denotes set cardinality.

| Approach | Max Input | Mean Input | $X \geq 12.5$ | $X \geq 13.0$ |
|----------|-----------|------------------------|---------------|---------------|
| EKG-MPPI | 1.037 | 1.187×10^{-4} | 3215 | 1176 |
| EnKC | 1.311 | 1.638×10^{-4} | 3293 | 1308 |
| MPPI | 1.780 | 1.547×10^{-4} | 3528 | 1284 |

325 Next, we investigate how changes in the hyperparameters of EKG-MPPI affect control performance. The temperature parameter λ is varied over $\{10, 1, 0.1, 0.01, 0.001\}$, and the number of samples is varied over $\{5, 10, 20, 40, 80\}$. For each configuration, we record the number of threshold exceedances, the average magnitude of the applied control input, and the number of control applications. The results are shown in Figure 3. As shown in Figure 3, decreasing λ generally increases the average control magnitude while decreasing the control frequency. In MPPI, λ controls how strongly the importance weights

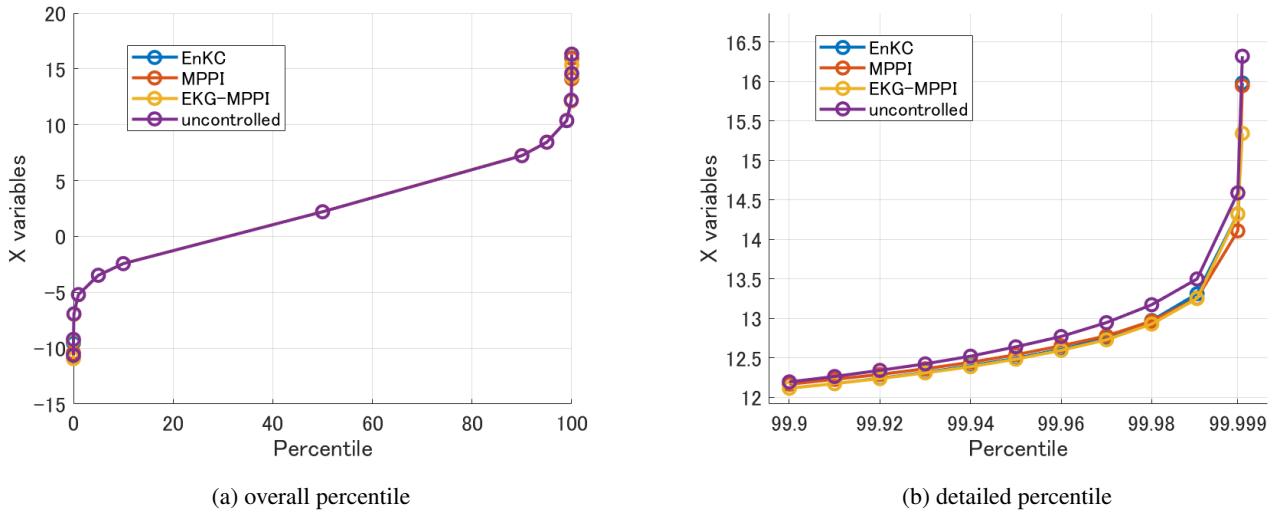


Figure 2. The percentile of state variables of the Lorenz 96 system in the control simulation experiments. While (a) shows the whole percentile ranges from 0 to 100, (b) shows the zoom-in of the range from 99.9 percentile to the maximum.

concentrate on low-cost samples: smaller λ yields more peaked weights and thus drives the update toward more aggressive perturbations found in the sampled rollouts. In our EKG-MPPI setting, the proposal distribution is EnKC-informed (with mean and variance depending on u_t^{EnKC}), so a smaller λ corresponds to a stronger deviation from this EnKC-guided proposal toward

330 the best-performing samples. When λ is very small (e.g., 0.01 or 0.001), increasing the sample size tends to further increase the average control magnitude and reduce the control frequency. A likely explanation is that, with more samples, the algorithm more often discovers rare, highly effective perturbations; when λ is small, the update places disproportionately large weight 335 on these samples, amplifying the resulting control magnitude. Consistently, Figure 3 also shows that larger sample sizes and smaller λ tend to suppress extreme events more strongly. An interesting non-monotonic behavior is observed for the average input: the smallest average input occurs around $\lambda = 1$ in this experiment. Although one might expect $\lambda = 10$ to yield smaller inputs, a too-large λ produces overly diffuse weights and thus overly conservative updates, which may allow extremes to persist and require larger (or more sustained) inputs later to suppress them.

Finally, we compare the computational time required to compute the control input at each control step. Table 2 summarizes the mean and standard deviation of the wall-clock time for calculating u_t^* over the simulation. As shown in the table, EKG-MPPI has the largest computational cost, followed by MPPI, while EnKC is the fastest. This ordering is expected: EKG-MPPI 340 combines the EnKC computation (ensemble forecasts and an EnKS-based update) with the MPPI refinement stage (multiple nonlinear rollouts), and is therefore more expensive than either component alone. The relatively high cost of MPPI in our implementation is mainly due to the repeated forward simulations required for each sample over the prediction horizon (i.e., $\mathcal{O}(K_{\text{MPPI}} T_{\text{MPPI}})$ model evaluations in a straightforward implementation). On the other hand, MPPI rollouts and weight 345 computations are embarrassingly parallel across samples, and can be accelerated substantially using multi-core CPUs and/or

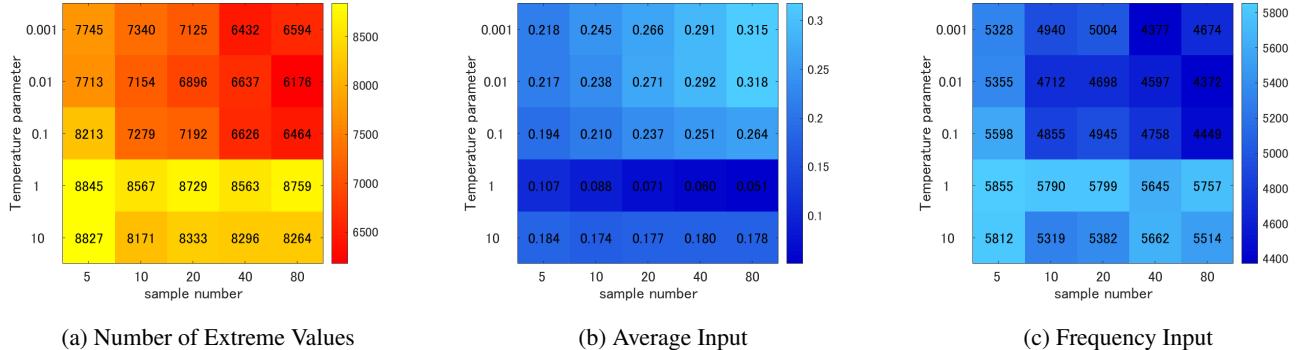


Figure 3. Hyperparameter configuration map displaying the effects of varying λ and sample size on control outcomes.

GPUs. Such acceleration directly benefits the MPPI stage within EKG-MPPI as well, and is therefore a promising route to reduce the overall wall-clock cost of the proposed method.

Table 2. Mean and Standard Deviation of the Control Input Computation Time (in sec) During Simulation

| Approach | Mean | Std |
|----------|-----------------------|-----------------------|
| EKG-MPPI | 9.21×10^{-4} | 1.12×10^{-3} |
| EnKC | 3.34×10^{-4} | 1.14×10^{-4} |
| MPPI | 7.77×10^{-4} | 1.96×10^{-4} |

4.2 Surface quasi-geostrophic model

The quasi-geostrophic (QG) model is a standard framework for describing mesoscale barotropic and baroclinic dynamics; for 350 a comprehensive review, see Vallis (2017). The surface quasi-geostrophic (SQG) model is derived from the QG model under the assumption of uniform interior potential vorticity. A key feature of the SQG model is that the surface buoyancy acts as an active tracer from which the horizontal velocity field is diagnosed. All numerical experiments follow the setup and parameter values in Resseguier et al. (2017).

4.2.1 SQG dynamics and diagnostic velocity

355 Let $b(x, t)$ denote the surface buoyancy on a doubly periodic domain $x = (x, y) \in \mathbb{T}^2$. In the SQG model, b evolves as an active tracer and the velocity is diagnosed from b via a streamfunction ψ :

$$\partial_t b + v \cdot \nabla b = \mathcal{D}(b) + f, \quad (41)$$

$$v = \nabla^\perp \psi = (-\partial_y \psi, \partial_x \psi), \quad (42)$$

$$b = (-\Delta)^{1/2} \psi. \quad (43)$$



360 Here $\mathcal{D}(b)$ denotes a dissipation operator and f is an external forcing term. Equivalently, in Fourier space, $\hat{\psi}(k) = |k|^{-1}\hat{b}(k)$ for $k \neq 0$, with the mean mode set to zero; see Resseguier et al. (2017) for numerical implementation details. We discretize b on a 128×128 grid and denote the vectorized buoyancy field at discrete time t by $x_t \in \mathbb{R}^n$ with $n = 128^2$. Let $M_0 : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be the one-step forecast model corresponding to the numerical time integrator of (41) (including the SQG inversion (43)):

$$x_{t+1} = M_0(x_t). \quad (44)$$

365 It is assumed that the surface buoyancy can be controlled via a small, spatially localized increment. Accordingly, we model the actuation as an additive perturbation applied at the control time t :

$$x_t^+ = x_t + u_t, \quad (45)$$

followed by the uncontrolled model integration. Thus, the controlled forecast model M used in Step 4 (Sect. 3.2) is defined by

$$M(x_t, u_t) = M_0(x_t + u_t), \quad M(x_t, 0) = M_0(x_t). \quad (46)$$

370 This is consistent with the single-step actuation setting assumed in Step 4 (Sect. 3.2), where the perturbation is applied only at the first rollout step and no further input is provided during the remaining rollout. We consider a single localized actuator at each control time. Let $l_t = (i_t, j_t)$ denote the actuator location on the 128×128 grid and $m_t \in \mathbb{R}$ its magnitude. Let $e_{l_t} \in \mathbb{R}^n$ be the standard basis vector corresponding to grid point l_t (i.e., a Kronecker delta on the vectorized grid). The control input is then

$$375 \quad u_t = m_t e_{l_t}. \quad (47)$$

In the EKG-MPPI/MPPI sampling step, l_t is sampled in \mathbb{R}^2 and then rounded to the nearest integer grid index and clipped to the valid range in each coordinate, following the same discretization convention as in Sect. 3.2. Given a buoyancy state x_t , we diagnose the streamfunction and velocity using (42)–(43), and define the wind-speed magnitude field as

$$w(x_t)(x) = \|v(x_t)(x)\|_2. \quad (48)$$

380 Let Ω_{tar} denote a prescribed target region (a set of grid points). Our control objective is to suppress wind speed within Ω_{tar} , and we define the running state-cost function as

$$S(x_t) = \text{target}(w(x_t)), \quad (49)$$

where $\text{target}(\cdot)$ aggregates wind speed within Ω_{tar} (e.g., regional mean or maximum). See Resseguier et al. (2017) for numerical details of the SQG inversion and the diagnostic computation of $w(\cdot)$.

385 **4.2.2 Simulation Results**

To compare the control performance of EKG-MPPI and EnKC, we conduct control simulations on a 128×128 grid map under eight different target regions: $[85, 95] \times [85, 95]$, $[85, 95] \times [90, 100]$, $[85, 95] \times [95, 105]$, $[90, 100] \times [85, 95]$, $[90, 100] \times$



[90, 100], [90, 100] \times [95, 105], [95, 105] \times [85, 95] and [95, 105] \times [90, 100]. Hereafter, these experiments are referred to as exp1 through exp8. One simulation is performed for each target region, resulting in a total of eight simulations. The total number 390 of simulation steps is set to 10118, which corresponds to approximately 10 days of real time. For EnKC, the ensemble size is set to $N = 40$, the prediction horizon is set to $T_{\text{EnKC}} = 500$ steps (approximately 12 hours), and the control weight is set to 10. For EKG-MPPI, the number of samples is set to $K_{\text{MPPI}} = 40$, and the prediction horizon is also set to 500 steps. The embedding function into the prior distribution is defined as follows:

$$\begin{aligned} f_{\text{loc}}(u_t^{\text{EnKC}}) &= \text{argnz}(u_t^{\text{EnKC}}), & g_{\text{loc}}(u_t^{\text{EnKC}}) &= 10.0, \\ f_{\text{mag}}(u_t^{\text{EnKC}}) &= u_t^{\text{EnKC}}(f_{\text{loc}}(u_t^{\text{EnKC}})), & g_{\text{mag}}(u_t^{\text{EnKC}}) &= \|u_t^{\text{EnKC}}\|_1, \end{aligned} \quad (50)$$

395 The state cost function for EKG-MPPI is defined as follows:

$$S(x_t) = \max(\text{target}(w(x_t)) - w_{th}, 0) \quad (51)$$

where $w(\cdot)$ is a function that converts the buoyancy state x_t into the corresponding wind speed, and $\text{target}(\cdot)$ is a function that computes the magnitude of the wind speed within the target region. Here w_{th} denotes the wind speed threshold, and in this 400 study we set $w_{th} = 2$. For details of the function $w(\cdot)$, the reader is referred to Resseguiet et al. (2017). The results are shown in Figure 4. EKG-MPPI achieves lower maximum wind speeds in the target region than EnKC (Figures 4a), while requiring 405 smaller mean control magnitudes (Figures 4b). One reason why EKG-MPPI demonstrates superior control performance compared to EnKC is that EKG-MPPI leverages prior information provided by EnKC to perform MPPI with high sample efficiency, enabling the computation of control inputs that explicitly account for the nonlinear dynamics of the atmospheric system. In Scenario 7, EKG-MPPI results in larger control inputs than EnKC. A possible explanation is that the difficulty of the control 410 task varies across target regions. In particular, the target region in Scenario 7, [85, 95] \times [95, 105], is characterized by strong wind speeds from the beginning of the simulation. In such regions, achieving the desired control performance is challenging even for EKG-MPPI, which may obscure the performance difference between EKG-MPPI and EnKC. As potential improvements, developing control methods that operate more robustly than EKG-MPPI in complex environments, or initiating control actions before wind speeds intensify, could be considered.

410 For clarity, Figure 5 shows visualizations of the no-control, EKG-MPPI, and EnKC simulations for the target region [85, 95] \times [90, 100]. From Figures 5a, 5b and 5c, it can be observed that strong winds occur in the target region in the absence of control. From 5d, 5e and 5f, EnKC suppresses the wind speed in the target region compared to the no-control case; however, localized 415 regions of strong wind still remain. From 5g, 5h and 5i, EKG-MPPI achieves stronger wind-speed suppression in the target region than EnKC.

415 4.3 Conclusions

We developed EnKC-guided MPPI (EKG-MPPI), a hybrid control scheme that uses EnKC to obtain a sparse candidate perturbation and then refines it by MPPI using nonlinear forward rollouts. The EnKC output is used to shape the sampling distribution in MPPI, so that exploration is concentrated around plausible actuator locations and magnitudes instead of relying on uninformed sampling.

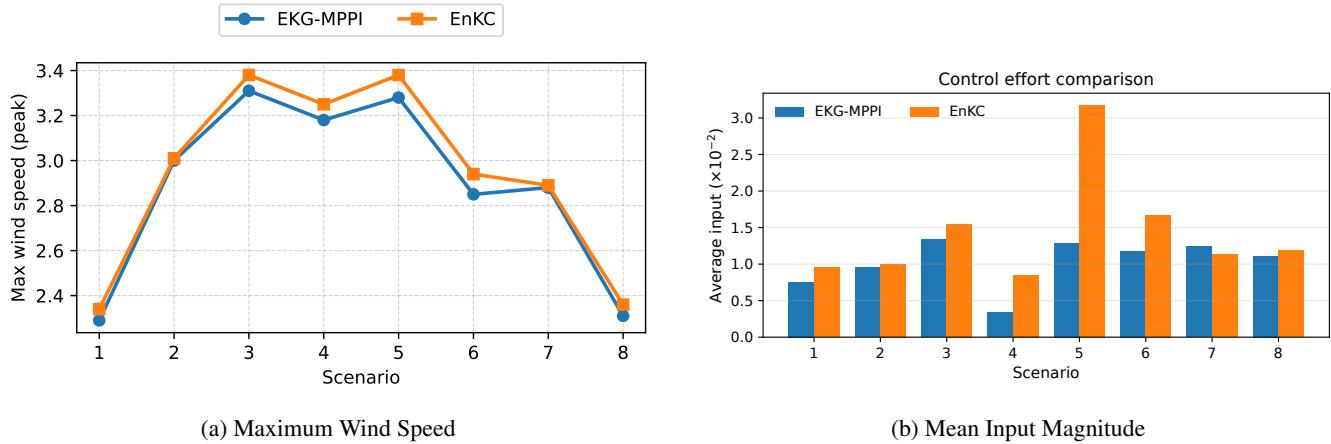


Figure 4. For each experiment, we compare the maximum wind speed and the mean input magnitude obtained by EKG-MPPI and EnKC. As an evaluation metric, we use the percentage difference defined as $(EKG - EnKC)/EnKC \times 100$. A negative value of this metric indicates that EKG-MPPI achieves a smaller maximum wind speed and mean input magnitude than EnKC, whereas a positive value indicates that EnKC yields smaller values.

420 In the Lorenz-96 control simulation experiment, EKG-MPPI reduced extreme-event counts compared with EnKC and vanilla MPPI, while requiring no larger (and often smaller) inputs. For example, the number of exceedances decreased from 3293 to 3215 at $X \geq 12.5$ and from 1308 to 1176 at $X \geq 13.0$ relative to EnKC, and the maximum input decreased from 1.311 to 1.037 (with the mean input decreasing from 1.638×10^{-4} to 1.187×10^{-4}). In the SQG experiments over eight target regions, EKG-MPPI also achieved lower maximum wind speeds in the target region than EnKC, while using smaller mean control magnitudes.

425 The current implementation assumes localized actuation (a single or very sparse actuator) and applies the perturbation as a one-step input within each MPPI rollout, and it does not yet impose hard physical constraints. Future work will address multi-actuator and multi-step actuation, constraint-aware formulations, and evaluation metrics that explicitly quantify non-target impacts and robustness to model/observation uncertainty, in addition to improving computational efficiency for higher-dimensional models.

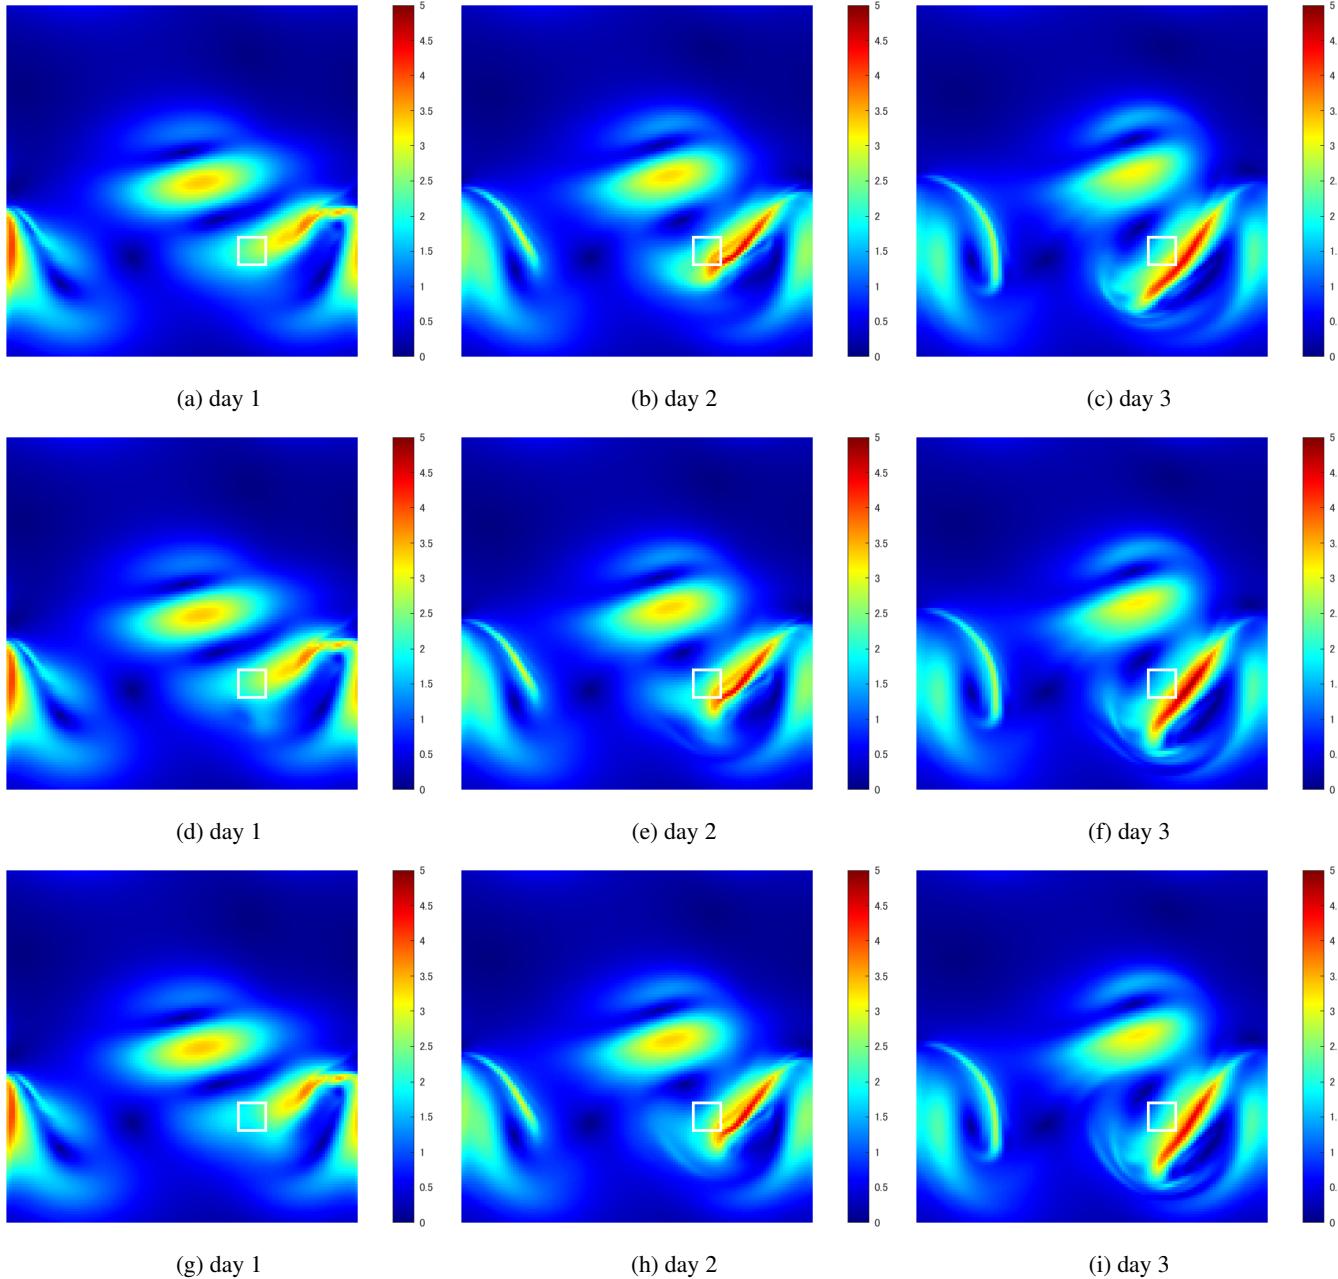


Figure 5. The first, second, and third rows show the time evolution of the wind-speed field for the no-control, EnKC, and EKG-MPPI cases, respectively. The white rectangle denotes the target region for wind-speed suppression.



References

Cotton, W., Zhang, H., Mcfarquhar, G., and Saleeby, S.: Should we consider polluting hurricanes to reduce their intensity, *J. Weather Mod.*, 39, 2007.

Houtekamer, P. L. and Zhang, F.: Review of the Ensemble Kalman Filter for Atmospheric Data Assimilation, *Monthly Weather Review*, 144, 4489 – 4532, <https://doi.org/10.1175/MWR-D-15-0440.1>, 2016.

435 Jacobson, M. and Kempton, W.: Taming hurricanes with arrays of offshore wind turbines, *Nature Climate Change*, 4, <https://doi.org/10.1038/nclimate2120>, 2014.

Kawasaki, F. and Kotsuki, S.: Leading the Lorenz 63 system toward the prescribed regime by model predictive control coupled with data assimilation, *Nonlinear Processes in Geophysics*, 31, 319–333, <https://doi.org/10.5194/npg-31-319-2024>, 2024.

440 Latham, J., Parkes, B., Gadian, A., and Salter, S.: Weakening of hurricanes via marine cloud brightening (MCB), *Atmospheric Science Letters*, 13, 231–237, <https://doi.org/10.1002/asl.402>, 2012.

Lorenz, E.: Predictability: a problem partly solved, Ph.D. thesis, Shinfield Park, Reading, 1995.

Miller, J., Tang, A., Tran, T. L., Prinsley, R., and Howden, M.: The Feasibility and Governance of Cyclone Interventions, *Climate Risk Management*, 41, 100535, <https://doi.org/10.1016/j.crm.2023.100535>, 2023.

445 Ouyang, M., Tokuda, K., and Kotsuki, S.: Reducing manipulations in a control simulation experiment based on instability vectors with the Lorenz-63 model, *Nonlinear Processes in Geophysics*, 30, 183–193, <https://doi.org/10.5194/npg-30-183-2023>, 2023.

Power, T. and Berenson, D.: Variational Inference MPC using Normalizing Flows and Out-of-Distribution Projection, <https://arxiv.org/abs/2205.04667>, 2022.

450 Resseguier, V., Mémin, E., and Chapron, B.: Geophysical flows under location uncertainty, Part III SQG and frontal dynamics under strong turbulence conditions, *Geophysical & Astrophysical Fluid Dynamics*, 111, 209–227, <https://doi.org/10.1080/03091929.2017.1312102>, 2017.

Sawada, Y.: Ensemble Kalman filter meets model predictive control in chaotic systems, <https://arxiv.org/abs/2403.06371>, 2024a.

Sawada, Y.: Quest for an efficient mathematical and computational method to explore optimal extreme weather modification, <https://arxiv.org/abs/2405.08387>, 2024b.

455 Schneider, T., Stuart, A. M., and Wu, J.-L.: Ensemble Kalman inversion for sparse learning of dynamical systems from time-averaged data, *Journal of Computational Physics*, 470, 111 559, <https://doi.org/10.1016/j.jcp.2022.111559>, 2022.

Sun, Q., Miyoshi, T., and Richard, S.: Control simulation experiments of extreme events with the Lorenz-96 model, *Nonlinear Processes in Geophysics*, 30, 117–128, <https://doi.org/10.5194/npg-30-117-2023>, 2023.

Vallis, G. K.: *Atmospheric and Oceanic Fluid Dynamics: Fundamentals and Large-Scale Circulation*, Cambridge University Press, 2 edn., 460 2017.

Williams, G., Drews, P., Goldfain, B., Rehg, J. M., and Theodorou, E. A.: Information-Theoretic Model Predictive Control: Theory and Applications to Autonomous Driving, *IEEE Transactions on Robotics*, 34, 1603–1622, <https://doi.org/10.1109/TRO.2018.2865891>, 2018.

Willoughby, H. E., Jorgensen, D. P., Black, R. A., and Rosenthal, S. L.: Project STORMFURY: A Scientific Chronicle 1962–1983, *Bulletin of the American Meteorological Society*, 66, 505 – 514, [https://doi.org/10.1175/1520-0477\(1985\)066<0505:PSASC>2.0.CO;2](https://doi.org/10.1175/1520-0477(1985)066<0505:PSASC>2.0.CO;2), 1985.

465 Zhang, H., Mcfarquhar, G., Saleeby, S., and Cotton, W.: Impacts of Dust in the SAL as CCN on the Evolution of an Idealized Tropical Cyclone, *AGU Fall Meeting Abstracts*, 2006.



Author contributions. **Haru Kuroki:** Conceptualization; investigation; methodology; validation; visualization; writing—original draft. **Kazumune Hashimoto:** Investigation; methodology; supervision; funding acquisition. **Yohei Sawada:** Investigation; methodology; supervision; funding acquisition. **Le Duc:** Investigation; methodology; **Masashi Minamide:** Investigation; writing—review and editing.

470 *Competing interests.* The authors declare no conflicts of interest.

Acknowledgements. This work was supported in part by the JST Moonshot R&D program under Grant JMPJMS2281.