

Supplementary Material

Additional comparison with the BO-informed MPPI baseline

This supplementary material provides additional details on the BO-informed MPPI baseline introduced in our response to Referee Comment 2 and summarizes the corresponding Lorenz–96 results. Throughout this document, we refer to this baseline as **BO-MPPI**.

Experimental setup

For each fixed value of the MPPI temperature parameter λ , BO-MPPI is run with

$$\alpha \in \{100, 10, 1, 0.1, 0.01\},$$

where α is the BO-side trade-off parameter that balances the predicted threshold exceedance over the rollout horizon against the intervention magnitude (see below). On the other hand, λ is the MPPI temperature parameter. The two parameters therefore play different roles and are varied separately.

To ensure a fair comparison with EKG-MPPI, BO-MPPI uses the same value of λ and the same sampling variances for intervention location and intervention magnitude as EKG-MPPI. In the BO stage, the *Expected Improvement* (EI) acquisition function is used, and the BO budget is fixed at 7 iterations at each control decision. This budget was selected because even 7 BO iterations already lead to a substantially larger per-decision computational cost than that of the proposed method.

The objective function used in Bayesian optimization evaluates both the magnitude of extreme events and the intervention cost. Let x_t denote the system state at time step t , u_t the intervention applied at time t , and r_{th} the threshold for extreme events. The system dynamics are represented by $\text{Model}(\cdot)$, and the prediction horizon is denoted by H . The cost function is defined as

$$c(x_t, u_t) = \sum_{i=1}^H \sum_{j=1}^N \max(x_{t+i}^{(j)} - r_{\text{th}}, 0) + \alpha \|u_t\|, \quad (\text{S1})$$

where the state trajectory is obtained by propagating the model dynamics as

$$x_{t+1} = \text{Model}(x_t + u_t), \quad x_{t+i} = \text{Model}(x_{t+i-1}) \quad (i = 2, \dots, H). \quad (\text{S2})$$

The first term measures the total exceedance above the threshold during the prediction horizon, while the second term penalizes the magnitude of the intervention. The parameter α controls the trade-off between intervention effectiveness and intervention magnitude.

Rescaling used for visualization

To facilitate visual comparison in the two-dimensional performance–control plots, we rescale both the average intervention magnitude and the threshold-exceedance count (the horizontal-axis performance metric) by the corresponding mean over the five BO-MPPI runs for the same fixed value of λ . This rescaling is used only for visualization. Because each quantity is divided by a

reference mean rather than mapped to the interval $[0, 1]$, we refer to the resulting quantities as *rescaled* (or *mean-scaled*) values rather than normalized values. All qualitative comparisons reported in the text are unchanged when the original, unscaled metrics are used.

For each fixed λ , let $\mathcal{A} = \{100, 10, 1, 0.1, 0.01\}$. Let m_{EKG} and m_α denote the average intervention magnitudes obtained by EKG-MPPI and BO-MPPI, respectively. We define the BO-MPPI reference mean by

$$\bar{m}_{\text{BO}} = \frac{1}{5} \sum_{\alpha \in \mathcal{A}} m_\alpha, \quad (\text{S3})$$

and the corresponding rescaled quantities by

$$\tilde{m}_{\text{EKG}} = \frac{m_{\text{EKG}}}{\bar{m}_{\text{BO}}}, \quad \tilde{m}_\alpha = \frac{m_\alpha}{\bar{m}_{\text{BO}}}. \quad (\text{S4})$$

Similarly, let e_{EKG} and e_α denote the threshold-exceedance counts used on the horizontal axis. We define

$$\bar{e}_{\text{BO}} = \frac{1}{5} \sum_{\alpha \in \mathcal{A}} e_\alpha, \quad (\text{S5})$$

and

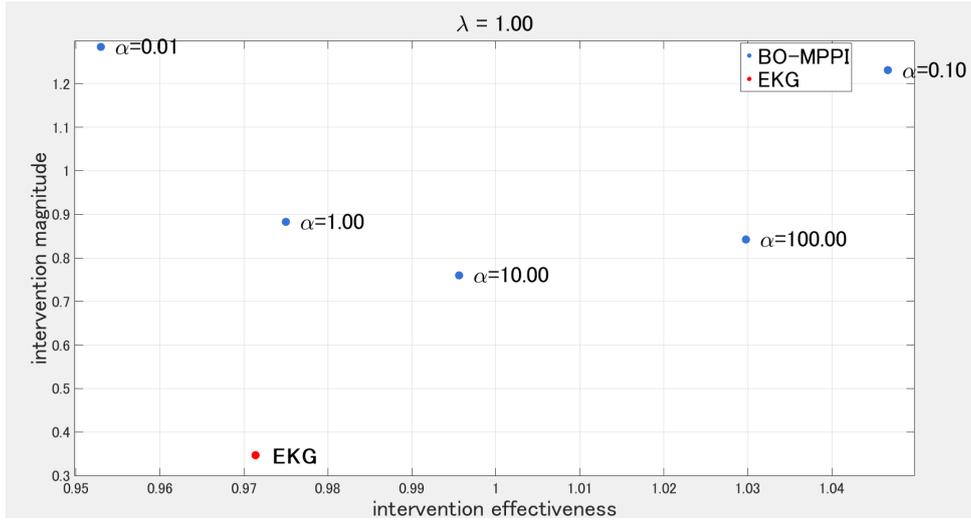
$$\tilde{e}_{\text{EKG}} = \frac{e_{\text{EKG}}}{\bar{e}_{\text{BO}}}, \quad \tilde{e}_\alpha = \frac{e_\alpha}{\bar{e}_{\text{BO}}}. \quad (\text{S6})$$

Values below 1 therefore indicates performance better than the BO-MPPI mean for the same λ on the corresponding axis. Smaller horizontal values indicate fewer threshold exceedances, and smaller vertical values indicate a smaller average intervention magnitude. Since each axis is rescaled by a common positive constant within each panel, this transformation does not change the within-panel ordering or Pareto dominance relations; it only changes the axis units for visualization.

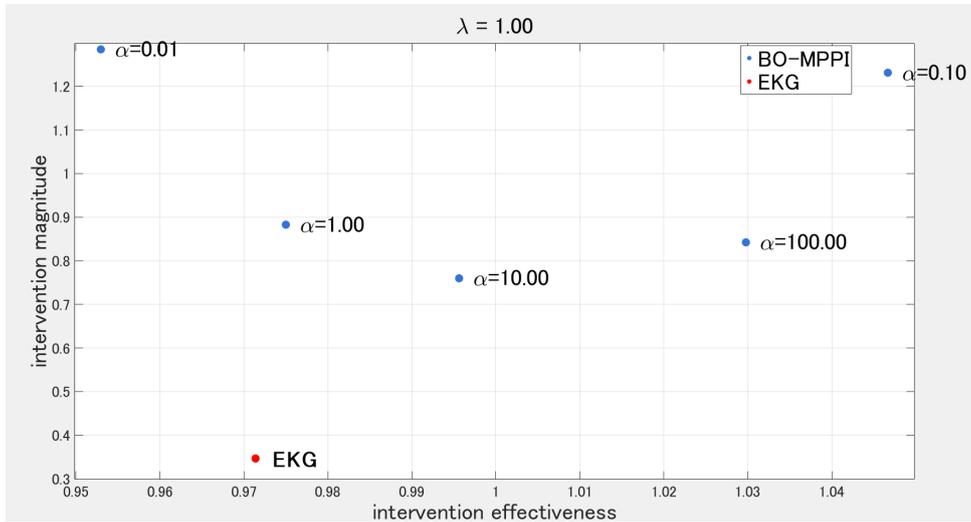
Results

Figure S1 compares EKG-MPPI and BO-MPPI for $\lambda \in \{0.1, 1.0, 10.0\}$. For each tested value of λ , the EKG-MPPI result lies below the empirical lower envelope (Pareto frontier-like trade-off curve) of the BO-MPPI results obtained by varying α . In other words, EKG-MPPI achieves a more favorable trade-off between the intervention-effectiveness metric on the horizontal axis and the average intervention magnitude on the vertical axis. Although EKG-MPPI is not uniformly located in the lower-left of all BO-MPPI points, no tested BO-MPPI configuration achieves both a more favorable intervention-effectiveness value and a smaller intervention magnitude than EKG-MPPI. This is consistent with the conclusion in our response to Referee Comment 2 that the advantage of EKG-MPPI is robust to the choice of λ and is not attributable to a single hyperparameter setting.

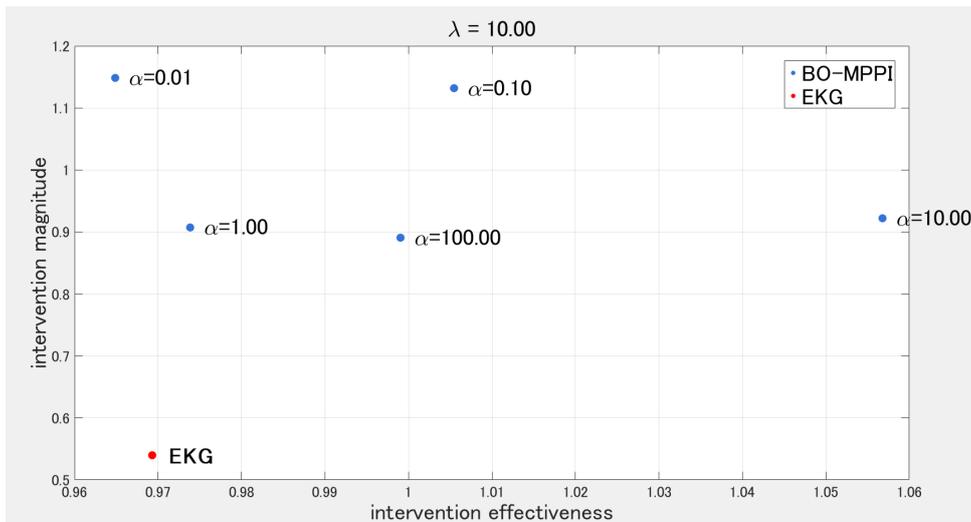
The mean computational time per decision step was 5.9953×10^{-2} s for BO-MPPI (variance: 2.3513×10^{-5} s²) and 2.80×10^{-4} s for EKG-MPPI (variance: 3.4×10^{-9} s²). Note that, since the sampling interval is 0.05 s, BO-MPPI slightly exceeds the available real-time budget on average, whereas EKG-MPPI remains well within this constraint.



(a) $\lambda = 0.1$



(b) $\lambda = 1.0$



(c) $\lambda = 10.0$

Figure S1: Comparison between EKG-MPPI and BO-MPPI in the Lorenz-96 experiment for three values of the MPPI temperature parameter λ . The horizontal axis represents the normalized intervention-effectiveness metric, defined here as the number of threshold exceedances, and the vertical axis represents the normalized average intervention magnitude. Blue points correspond to BO-MPPI with different values of the BO trade-off parameter α , and the red point corresponds to EKG-MPPI.