

The paper pushes weather intervention from long-term, open-loop, climate-scale schemes toward closed-loop control of individual rainfall events, and the methodological chain (MPC + gradient-free sampling + physics-informed prior) is coherently designed and clearly written. However, the support for the central scientific claims is not yet sufficient: the mechanistic interpretation relies on chaotic perturbation growth without any test against model error; the headline advantage over random placement is only marginally significant and lacks an ablation; the redistribution of precipitation (the side effect) is not quantified; and only a single event is tested. Code and data are offered only “upon reasonable request”.

Major comments

MS1. The hypothesized mechanism relies on chaotic perturbation growth, and the identical-twin setup cannot demonstrate real-world viability.

The authors themselves state (p. 13): “The reduction in target-region precipitation cannot be directly explained by local changes in the effective wind and humidity fields and therefore cannot be attributed solely to local drying or weakened moisture transport. Instead, the results suggest a plausible mechanism in which the distributed actuators may stimulate mesoscale perturbation growth in the atmosphere (Ono and Inatsu, 2025), thereby altering the atmospheric dynamics and reorganizing the wind, humidity, and moisture transport fields.” If the rainfall reduction is achieved mainly through chaotic trajectory divergence rather than through a deterministic momentum-extraction effect, then (i) in a real forecast with model and initial-condition errors, the sign and magnitude of the response to the same intervention become unpredictable, and the intervention could just as well enhance rainfall; and (ii) all experiments in this paper use the same model for prediction and for the “real” system (only near-surface moisture is perturbed), which is a classical identical-twin configuration that masks exactly this problem. I therefore ask the authors to: (a) discuss this limitation explicitly; (b) add sensitivity tests against model error (e.g., a different cloud-microphysics scheme, stochastic kinetic-energy-backscatter-type perturbations, or an assessment of whether the sign of the precipitation response is robust across different perturbation realizations); and (c) provide a physical argument for why the perturbation growth triggered by the sails should preferentially act to reduce rainfall over the target region, rather than citing the mesoscale perturbation-growth literature alone.

MS2. The core conclusion that the physics-informed prior outperforms random placement is statistically and experimentally under-supported.

(a) With only 10 ensemble members, the improvement of the humidity-informed mean over random placement (23.36 mm) is only marginally significant: my own paired t-test on Table 2 gives $t \approx 2.3$ (two-tailed $p \approx 0.047$), and a Wilcoxon signed-rank test sits at the critical value. Moreover, the sample standard deviation of the humidity-informed ensemble (20.97 mm, p. 14) is roughly twice that of the uncontrolled ensemble (≈ 9.3 mm, computed from Table 2), so in an absolute sense the outcomes are not more consistent; the statement “showing that the proposed strategy yields more consistent outcomes” (p. 14) holds only relative to random placement, and the uncontrolled ensemble standard deviation is not reported. (b) The random-placement baseline randomizes both actuator centers and orientation angles — “actuator centers and orientation angles are uniformly sampled within the control region and from $[-\pi, \pi)$, respectively” (p. 12) —

whereas the humidity-informed strategy additionally uses wind-aligned orientations via Eq. (12). The present comparison therefore cannot attribute the improvement to the humidity-informed CVT placement itself rather than to orientation alignment. Ablation experiments are needed (e.g., random locations + wind-aligned orientations; CVT locations + random orientations), together with statistical tests and effect sizes for all comparisons. (c) The objective in Eq. (6) is an expected cost, but the implementation uses $M = 1$ (single Monte Carlo roll-out, Table 1) when selecting the candidate with "the minimum predicted cost" via Eq. (8); selection based on a single noisy realization is at odds with the claimed robustness. Please discuss the sensitivity to M or adopt a more robust selection criterion.

MS3. The redistribution of precipitation (side effects) is not quantified and should not be deferred entirely to future work.

Figure 2(d) clearly shows that the rainfall reduction over the target region is accompanied by precipitation increases in several adjacent regions ("accumulated precipitation decreases substantially over the target region but increases in several adjacent regions", p. 13). For any rainfall-intervention method, the key question is whether it merely displaces flood risk to neighboring or downstream areas. This should be quantified at least to the extent of: the location and 72-hour accumulated rainfall increase of the affected regions, the ratio of the additional rainfall there to the rainfall reduction over the target region, and whether the additional rainfall falls over land, populated areas, or major shipping lanes. The promise in the conclusions to incorporate "spatial risk measures and constraints" should be backed by a preliminary quantitative analysis in this paper.

Minor comments

1. A single event (Kyushu, 10–13 August 2021) is insufficient to support the claim that "its general formulation allows it to be extended to other weather systems" (p. 16). Please add at least one contrasting case (e.g., a typhoon or purely local convective event), or clearly qualify the generality statement.
2. The uncertainty representation is narrow: "an additive random perturbation is applied at each control step to the water vapor mixing ratio throughout the horizontal computational domain within the 0–300 m layer", "uniformly sampled within $\pm 0.001 \text{ kg} \cdot \text{kg}^{-1}$ " (p. 11). Real forecast uncertainty extends through the tropospheric column and involves multiple variables; because this perturbation coincides with the 300 m layer in which the sails act, it may artificially inflate the apparent controllability. Please provide sensitivity tests to the perturbation magnitude, the perturbed variable(s), and the vertical extent.
3. The physical parameters of the sails (size, drag coefficient, the form of the additional drag forcing, and the grid-scale conversion of $d = 4.8 \text{ km}$) are not given, so the experiments cannot be reproduced; the assumption that "Actuator relocation and reorientation are assumed to occur instantaneously at each control update" (p. 11) also deserves a discussion of engineering feasibility (a 15 km displacement within a 12-hour interval).
4. The real-time argument — "takes approximately 7 hours", so that "the computational time remains shorter than the corresponding physical control interval, supporting the feasibility of real-time closed-loop precipitation regulation" (p. 11) — relies on the Fugaku supercomputer;

please comment on the resource requirements for practical deployment and temper the wording given that these are offline case-study experiments.

5. Eq. (11): if the wind speed at all levels is zero at a location, the weights sum to zero, $\mathbf{v}_{eff}$ is the zero vector, and the orientation from Eq. (12) is undefined; please state the numerical remedy.
6. Notational inconsistency: Eq. (14) defines $H_t(P_t)$, but the text reads “minimizing the locational cost $H(P)$ yields a centroidal Voronoi tessellation” (p. 9); please unify.
7. Fig. 4 shows the layouts of a single case only; please add, for each control step, the predicted cost J of the selected layout relative to the candidate set, to visualize the sampling–selection process.
8. “The code and data supporting the findings of this study are available from the corresponding author upon reasonable request” (p. 16) does not meet current Copernicus data-policy expectations; please deposit code and data in a DOI-issuing repository.
9. As a method that deliberately alters regional precipitation, the paper would benefit from a short discussion of societal, ethical, and cross-border regulatory implications.

Specific language/spelling/formatting issues (with verbatim quotes)

10. p. 6, first paragraph of Sect. 3: “...while decomposing the long-term regulation task into tractable finite-horizon layout optimization problem” → should be the plural “problems”(or “a ... problem”).
11. p. 14: “an additional decrease of 23.36 mm compared with random-placement ensemble” → “compared with the random-placement ensemble”.
12. p. 9: “minimizing the locational cost $H(P)$ ” → Eq. (14) defines $H_t(P_t)$; please align.
13. Diamond et al., 2022: “119, e2118379 119” — the volume number “119” is duplicated; it should read “119, e2118379”.
14. Feingold et al., 2024: the reference list uses “et al.” without listing all authors, inconsistent with the other entries (Copernicus style requires full author lists or a consistent convention).
15. Abstract vs. p. 3: “This is key to enhancing the robustness and practical viability of sampling-based optimization for such a complex rainfall control problem” (Abstract) is nearly identical to “This design is essential for enhancing the robustness and practical viability of sampling-based methods in such a complex rainfall control problem” (end of Sect. 1); please rephrase one of them.
16. p. 15: “Ensemble-based evaluations demonstrated the effectiveness and robustness...” → the paper otherwise uses the present tense; suggest “demonstrate”.
17. Table 1: “Parameter settings for the humidity-informed experiments” — the parameters in Table 1 are shared by all three ensembles, not only the humidity-informed one; suggest “for the numerical experiments”
18. p. 11: “corresponding to approximately 6% of the reference moisture value of $0.018 \text{ kg} \cdot \text{kg}^{-1}$ — $0.001/0.018 \approx 5.6\%$, and the perturbation range ± 0.001 has a full width of 0.002 ($\approx 11\%$); please state whether 6% refers to the half-width or the full width.
19. The time axis reads “1, 6, 12, 18, ...” — the first tick is 1 while the remaining ticks are spaced by 6 h; please use either “0, 6, 12, ...” or “1, 12, 24, ...”.