Physics-Informed Graph Diffusion for Climate Downscaling

Anonymous Submission

0.1 Motivation

High-resolution (≤10–25 km) surface fields are required for local risk analyses and extreme event attribution, yet orography and sparse stations limit standard downscaling. Diffusion models excel at synthesizing detail, but climate fields must obey basic laws: temperature decreases with elevation (environmental lapse rate), precipitation exhibits spatial coherence, and cross-variable structure should be physically reasonable. We ask: *Can generative models produce high-resolution fields that remain physically plausible, not just visually detailed?*

0.2 Method

Graph representation. We discretize the region into grid cells (nodes) with edges connecting k-nearest neighbors by geodesic distance and optional topographic similarity. Node features include annualized covariates (e.g., temperature and precipitation anomalies, nightlights); edges carry distances/slopes when available.

Diffusion on graphs. A transformer U-Net operates on node features with message-passing blocks; the denoising score function is conditioned on (i) diffusion timestep embeddings and (ii) graph context (node/edge features, pooled summaries).

Physics losses (differentiable). We add light regularizers to the diffusion objective:

- *Lapse-rate prior:* penalize positive temperature–elevation correlation ($\uparrow T$ with $\uparrow z$) and overly weak negative slopes; encourages realistic cooling with height.
- *Precipitation coherence:* graph-Laplacian penalty on precip anomalies to discourage speckle while allowing mesoscale structure.
- *Cross-variable consistency:* weak bounds on local temperature–precip correlation to avoid unphysical co-variability.

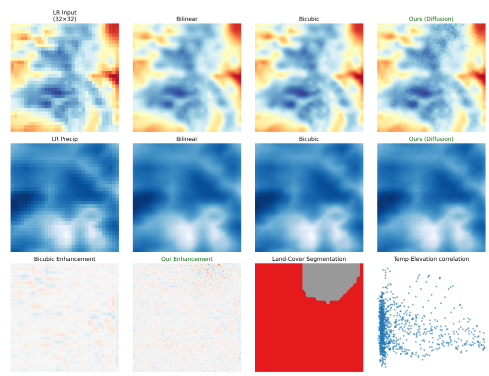

Figure 1: ERA5-Land (CO) qualitative comparison. Top: temperature (LR 32×32; 4× SR). Middle: precipitation. Bottom: enhancement diffs and temp–elevation correlation. Physics-informed diffusion preserves lapse-rate structure and precipitation coherence.

An auxiliary land-cover head (cross-entropy) provides semantic guidance; the total loss is a weighted sum of diffusion, physics, and auxiliary terms. Training uses AdamW with cosine scheduling.

Results. On ERA5-Land (Colorado, 2021–2022), bicubic interpolation achieves the lowest RMSE (0.130) and highest SSIM (0.882), while our physics-informed graph diffusion yields higher error (RMSE 0.210, SSIM 0.731) but substantially better physics compliance (0.68 vs 0.42–0.45 for interpolation).

0.3 Key Insights

(i) **Baselines matter:** classical interpolation remains a strong competitor on RMSE. (ii) **Physics vs. pixels:** enforcing physical constraints often increases numerical error, underscoring the need for evaluation metrics aligned with scientific end-use. (iii) **Lightweight priors work:** simple, differentiable physics regularizers can improve plausibility without expensive PDE solvers. (iv) **Future directions:** curriculum/homotopy schedules for balancing objectives, uncertainty calibration, and scaling to multi-year graph embeddings (e.g., AlphaEarth) for discovering long-term climate trends.