

WeatherEdit: Controllable Weather Editing with 4D Gaussian Field

Chenghao Qian¹, Wenjing Li^{1,2,†}, Yuhu Guo³, Gustav Markkula¹

¹University of Leeds

²Hefei University of Technology

³Carnegie Mellon University

tscq@leeds.ac.uk, W.Li@leeds.ac.uk, yuhu.guo@andrew.cmu.edu, G.Markkula@leeds.ac.uk

Abstract

In this work, we present WeatherEdit, a novel weather editing pipeline for generating realistic weather effects with controllable types and severity in 3D scenes. Our approach is structured into two key components: weather background editing and weather particle construction. For weather background editing, we introduce an all-in-one adapter that integrates multiple weather styles into a single diffusion model, enabling the generation of diverse weather effects in 2D image backgrounds. During inference, we design a Temporal-View (TV-) attention mechanism that follows a specific order to aggregate temporal and spatial information, ensuring consistent editing across multi-frame and multi-view images. To construct the weather particles, we first reconstruct a 3D scene using the edited images and then introduce a 4D Gaussian field to generate snowflakes, raindrops and fog in the scene. The attributes and dynamics of these particles are controlled through attribute modelling and dynamic simulation, ensuring realistic weather representation and flexible severity adjustments. Finally, we integrate the 4D Gaussian field with the 3D scene to render consistent and highly realistic weather effects. Experiments on multiple driving datasets demonstrate that WeatherEdit can generate diverse weather effects with controllable condition severity, highlighting its potential for autonomous driving simulation in adverse weather.

Code — <https://github.com/Jumponthemoon/WeatherEdit>

Introduction

3D weather editing (3D-WE) aims to generate multi-weather conditions in 3D scenes based on user instructions. It is a crucial task in computer vision with applications in autonomous driving (Qian et al. 2025c), augmented reality, and virtual scene synthesis (Almalioglu et al. 2022). While existing methods range from 2D-based image techniques (Valanarasu, Yasarla, and Patel 2022; Zhang et al. 2023) to 3D neural rendering (Li et al. 2023), developing a unified framework that enables multiple realistic weather effects and controllable severity remains a significant challenge. For instance, style transfer methods can modify weather backgrounds but are typically limited to a single effect.

[†] indicates corresponding author.

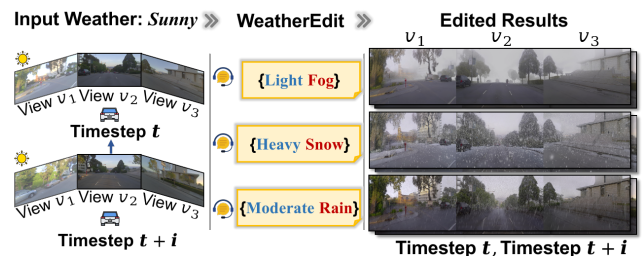


Figure 1: Realistic, Controllable Weather Editing. Given sunny-weather images captured by multi-view cameras over multiple timesteps, WeatherEdit can generate high-fidelity scenarios under diverse weather conditions (e.g., foggy, snowy, and rainy) with various severity levels (e.g., light, moderate, heavy) based on user instructions.

Diffusion models (Yang et al. 2024; Qian et al. 2025b) enable multi-weather synthesis; however, they often lack fine-grained control and struggle with spatial and temporal consistency across multiple images (Rothmeier and Huber 2021). More recently, ClimateNeRF (Li et al. 2023) utilized Neural Radiance Fields (NeRF) to model snow accumulation and flooding in 3D space. However, this method is limited to static weather events and leaves the overall weather tone unchanged, lacking full-scene atmospheric realism.

In 3D-WE, achieving both **realism** and **controllability** is important. Realism ensures high-quality weather effects and we consider it involves three aspects: (R1) *Distortion-free editing*: The editing results should be free from visual distortions and unnatural artifacts. (R2) *Temporal-view consistency*: Weather effects and scene elements should maintain consistency across both time steps and different viewpoints. (R3) *Dynamic weather effects*: Weather particles, such as snowflakes and raindrops, should be realistically simulated to reflect natural weather dynamics. On the other hand, controllability ensures flexible and fine-grained editing of weather effects, enabling users to customize various aspects of the scene, including: (C1) *Type*: The system should support the selection and transition between different weather types (e.g., rainy, snowy, foggy). (C2) *Event*: The system should allow users to specify whether the weather effect is static (e.g., previously fallen snow) or ongoing (e.g., falling

snow). (C3) *Severity*: Users should be able to adjust the intensity of weather conditions, such as changing light rain to heavy rain or increasing fog density. Unfortunately, current approaches have not explicitly considered the above aspects, leading to limitations in realism and controllability.

In this paper, we present WeatherEdit, a novel weather editing framework that progressively transforms 2D images into 4D dynamic weather scenes. This is achieved through a two-step process: **weather background editing** and **weather particle construction**.

In weather background editing, we first introduce an efficient all-in-one adapter that integrates multiple weather styles into a single model, allowing for the fine-tuning of a diffusion model to achieve distortion-free rendering (R1) and flexible weather type control (C1). Second, to ensure temporal-view consistency (R2), we introduce a parameter-free attention mechanism that aligns different viewpoints simultaneously and the same viewpoint over time. For weather particle construction, we propose a novel 4D Gaussian field, representing weather particles as 3D Gaussians with attributes such as color and density while simulating their motion. This enables dynamic weather effects (R3) and severity (C3). Rather than simulating an infinite field for full-scene coverage, we optimize efficiency by recycling particles within a compact field and aligning it with the rendering camera, maintaining both realism and performance.

The final weather effect is rendered from the 4D Gaussian field combined with the 3D scene reconstructed using edited weather background images. This flexible combination allows control over weather events, such as a wet road without rainfall (static) or a heavy snowstorm with limited visibility (dynamic), enabling adjustable event (C2). The main points of this paper are summarized below:

1. We propose WeatherEdit, a novel progressive 2D-to-4D framework that enables realistic and controllable dynamic weather editing.
2. We introduce an all-in-one adapter to enable a diffusion model for multi-weather (snowy, rainy and fog) synthesis and a Temporal-View attention to ensure consistent editing across multi-frame and multi-view.
3. We design a 4D Gaussian field for weather particle construction, enabling plausible simulation of raindrops, snowflakes, and fog with controllable severity.
4. We demonstrate WeatherEdit’s effectiveness in generating weather effects in 3D driving scenes, showcasing its applicability to real-world scenarios.

Related Work

Weather Effect Synthesis

Existing methods can be categorized into 2D-based and 3D-based approaches. 2D-based methods utilize generative models (Qian et al. 2025b; Parmar et al. 2024) to synthesize weather effects in single images. (Hwang et al. 2022; Schmidt et al. 2019) use CycleGAN (Zhu et al. 2017) to edit weather conditions, while (Schmidt et al. 2022) applies GAN-based inpainting to generate flooding with water reflections. (Hahner et al. 2019) uses a graphical model to

produce fog effects. Conventional 3D-based methods (Stomakhin et al. 2013; Zsolnai-Fehér 2022; Feldman and O’Brien 2002; Haas 2014) use graphical simulation to model weather in 3D space. More recently, (Li et al. 2023) has introduced neural rendering techniques for creating weather backgrounds but it is limited to post-weather event.

Diffusion Model for Image Editing

Diffusion models (Ho, Jain, and Abbeel 2020; Meng et al. 2021; Rombach et al. 2021) have significantly advanced image editing, enabling flexible and controllable modifications. Methods such as ControlNet (Zhang, Rao, and Agrawala 2023) and T2I-Adapter (Mou et al. 2024) introduce trainable modules to guide the generation process, allowing for fine-grained control over diffusion-based outputs. Instruct-Pix2Pix (Brooks, Holynski, and Efros 2023) enables interactive editing with textual instructions by trained on image-text pairs. TurboEdit (Deutch et al. 2024) employs a pseudo-guidance approach to enhance edit magnitude without introducing artifacts.

Although these methods support diverse styles, they are not specifically designed for weather editing, often distorting content or generating overly artistic effects. Moreover, they primarily focus on single-image editing, lacking spatial and temporal consistency across multiple frames and views. These limitations restrict their ability to generate realistic and consistent weather effects within the same scene.

3D Scene Reconstruction

Neural Radiance Fields (NeRF) (Mildenhall et al. 2020) and 3D Gaussian Splatting (3DGS) (Kerbl et al. 2023) are leading techniques for 3D scene reconstruction from multi-view images. NeRF encodes scenes implicitly with a neural network, while 3DGS represents them explicitly as 3D Gaussians, enabling faster rendering. Their versatility allows applications in dynamic scene reconstruction (Chen et al. 2025; Yan, Li, and Lee 2023), content creation (Liu et al. 2024; Lin et al. 2023), and artifact removal (Qian et al. 2025a). While (Li et al. 2023) extends NeRF to model smog, snow, and floods for weather simulation but focuses only on static weather conditions, leaving dynamic weather phenomena unaddressed.

Methods

In order to generate realistic and controllable multi-weather scenes in 3D, we propose a novel framework, WeatherEdit, which follows a progressive 2D-to-4D transformation process. As illustrated in Figure 2, WeatherEdit adopts a two-step weather editing pipeline: weather background editing and weather particle construction. The first step generates distortion-free, spatially and temporally aligned 2D images with edited backgrounds, while the second step dynamically models weather particles with various types and severity levels. The joint association of these two steps enables the flexible generation of multi-weather scenes.

Weather Background Editing

To support multiple editing styles, we develop a unified adapter to finetune the diffusion model and further enhance

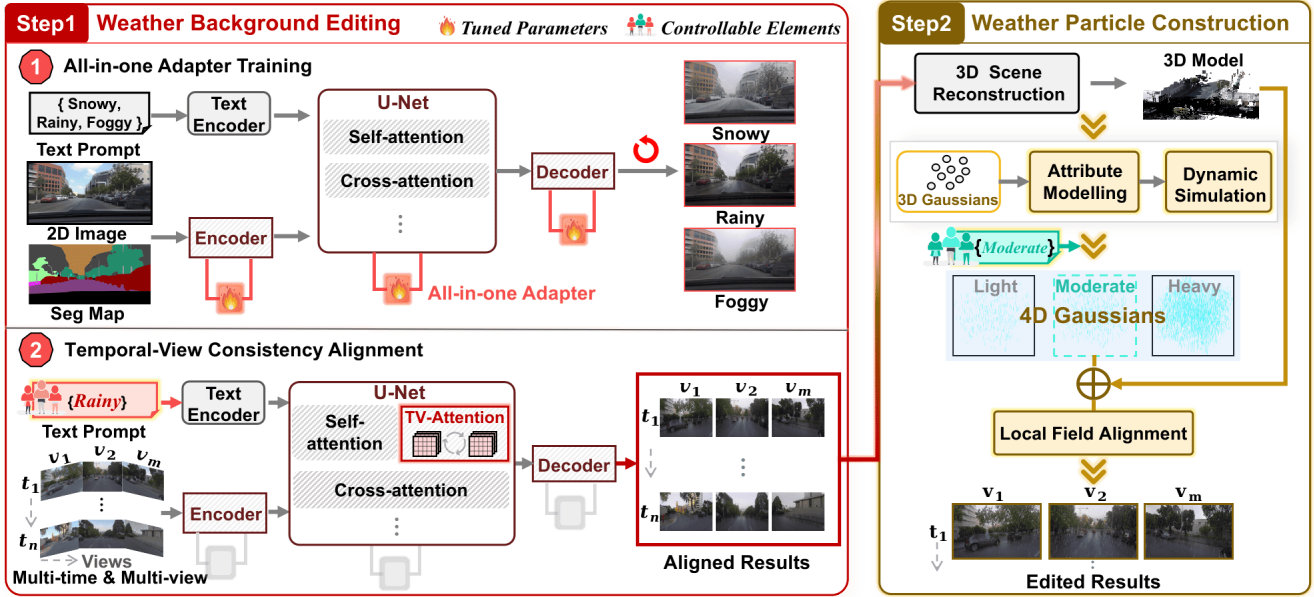


Figure 2: Overview of the WeatherEdit Framework. (a) Weather Background Editing: An all-in-one adapter enables the synthesis of diverse weather effects (e.g., snow, rain, fog) from text prompts and segmentation maps. Inference takes multi-frame and multi-view images as input, using temporal-view (TV) attention to ensure consistency. (b) Weather Particle Construction: A 4D Gaussian field, initialized from 3D Gaussians, undergoes attribute modeling and dynamic simulation, allowing user-controlled weather intensity (e.g., light, moderate, heavy). Local field alignment ensures seamless integration of the reconstructed 3D scene with dynamic weather effects.




Figure 3: Single-style Adapters vs. All-in-One Adapter. (a) Utilizes separate adapters for each weather style (e.g., snow, rain, fog), while (b) employs a unified all-in-one adapter for efficient multi-style adaptation.

the results through semantically conditioned inputs. Additionally, we introduce a novel Temporal-View Attention mechanism to ensure consistency in the edited results.

All-in-One Adapter. Existing diffusion-based methods often suffer from scene deformations and excessive stylization, making them unsuitable for precise and realistic weather editing. Although (Parmar et al. 2024) achieves effective style editing by fine-tuning the models with multiple single-style LoRA (Hu et al. 2022) adapters (shown in Figure 3 (a)), it requires training a separate model for each style, resulting in increased computational cost and limited flexibility.

To address this, we develop an all-in-one LoRA adapter that injects multiple weather styles into a single model (as shown in Figure 3 (b)), eliminating the need for training multiple ones. Let $W_0 \in R^{d \times k}$ be the pre-trained weights, x

the input image, and i the weather style (e.g., snowy, rainy, foggy) specified by text prompt. Each style is associated with a low-rank matrix L_i . The adapted model’s forward pass is then computed as $h = W_0x + L_i x$.

Moreover, since weather effects depend on scene semantics (e.g., snow on trees, rain on roads), we enhance editing by conditioning the model on semantic segmentation maps. Given an input image x with its corresponding segmentation map M , we encode the latent representation as $z_0 = \mathcal{E}(x, M)$. The model fine-tuning objective is then formulated as:

$$\mathcal{L}(\Delta\theta) = E_{\epsilon, t} [\|\epsilon - \epsilon_{\theta+\Delta\theta}(\sqrt{a_t}z_0 + \sqrt{1-a_t}\epsilon, t, p)\|^2], \quad (1)$$

where ϵ is the sampled Gaussian noise, $\epsilon_{\theta+\Delta\theta}$ is the UNet model with LoRA, and p is the text prompt. This formulation ensures editing results are semantically relevant, enabling precise and context-aware weather editing.

Temporal-View Consistency Alignment. After fine-tuning the model, multi-frame and multi-view images are processed to achieve background scene editing. However, processing each image independently can lead to inconsistencies across time steps and viewpoints due to the inherent stochasticity of generative models. To address this, we propose a Temporal-View (TV-) attention mechanism (as shown in Figure 4) to enable our model to integrate information from adjacent frames and key viewpoints.

In multi-view driving datasets, cameras are typically positioned in a structured left-front-right arrangement. The front

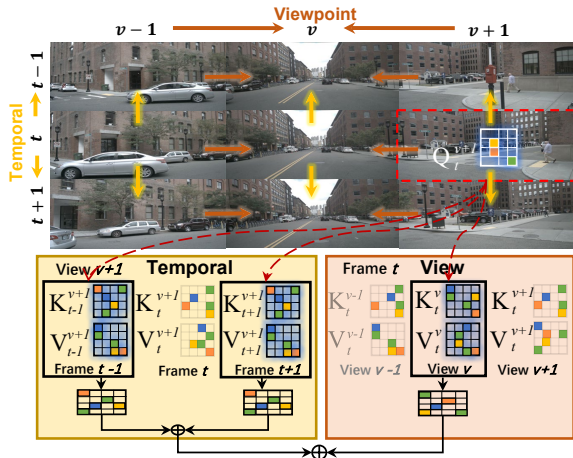


Figure 4: Temporal-View Attention for Consistent Weather Editing. The TV-attention aggregates information across both adjacent frames and center viewpoints to ensure multi-frame and multi-viewpoint consistency. The top section illustrates frame and viewpoint relationships, indicated by directional arrows, while the bottom section visualizes the attention operations.

camera captures overlapping content with both the left and right cameras, providing shared information across views. To leverage this, we design the view attention to allow the left and right cameras to query the center (front) view:

$$\text{S_Attn}_t^v = \text{Softmax} \left(\frac{Q_t^v \cdot K_t^{v,T}}{\sqrt{d}} \right) V_t^v, \quad (2)$$

where Q_t^v is the query from view v , and K_t^v, V_t^v are the key and value derived from the center view v . This ensures that each side viewpoint aggregates information from the center view, maintaining result consistency across different perspectives.

To preserve temporal coherence, each frame t queries its adjacent frames ($t-1$ and $t+1$), allowing it to integrate contextual information from both past and future frames. The temporal attention is defined as:

$$\text{T_Attn}_t^v = \text{Softmax} \left(\frac{Q_t^v \cdot K_v^{\mathcal{T}(t),T}}{\sqrt{d}} \right) V_v^{\mathcal{T}(t)}, \quad (3)$$

where $\mathcal{T}(t) = \{t-1, t+1\}$, ensuring that each frame incorporates temporal context, preventing inconsistencies across consecutive frames. The final attention is computed as a weighted sum of self-, view-, and temporal-attention, with λ controlling the balance to ensure global coherence across frames and viewpoints:

$$\text{Attn}_t^v = \lambda \cdot \text{SelfAttn}_t^v + (1 - \lambda) \cdot (\text{S_Attn}_t^v + \text{T_Attn}_t^v). \quad (4)$$

In this way, we can effectively reduce inconsistencies caused by independent image processing, producing a coherent editing effect across both view and temporal domains.

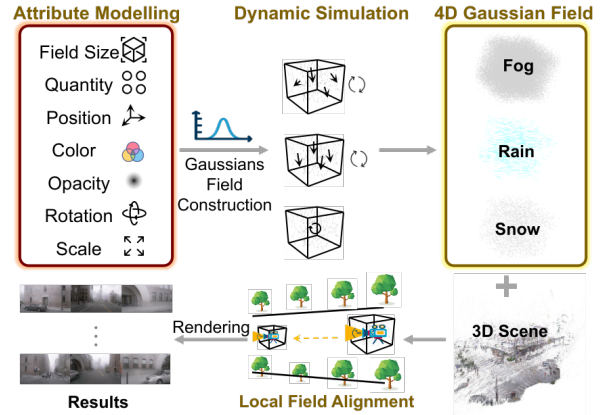


Figure 5: Weather Particle Construction Process. The attribute modeling stage defines weather properties to construct a 4D Gaussian field, which then undergoes dynamic simulation to model fog, rain, and snow. Local field alignment ensures seamless integration with the 3D scene, enabling realistic weather rendering.

Weather Particle Construction

Real-world weather is inherently regional, and we incorporate this into weather particle construction. Specifically, as shown in Figure 5, we propose a *4D Gaussian field*, initialized with 3D Gaussians carrying physical attributes (e.g., color, quantity) to model raindrops, fog, and snowflakes through attribute modeling. To enhance realism, we apply dynamic simulation to create realistic weather effects.

To generate a consistent falling particle effect, a large number of particles must be simulated, which is computationally expensive. To address this, we model particle behavior within a localized field, dynamically recycling particles as they exit while aligning the field with the rendering camera. This ensures a continuous and realistic weather effect while optimizing computational efficiency.

Attribute Modeling. Different weather particles exhibit unique characteristics. For instance, snowflakes are typically white with irregular shapes, while raindrops are transparent and elongated. To accurately capture these differences, we define a set of characteristic attributes A to distinguish different types of weather particles: $A_i = \{C_i, P_i, R_i, S_i, O_i\}$, where, C_i defines color, P_i represents position, R_i and S_i correspond to rotation and scale, and O_i signifies opacity. The index i refers to a specific type of weather, such as snow, rain, or fog. Additionally, we model these parameters using Gaussian distributions, expressed as $G(A_i)$, to mimic the natural randomness of real-world weather. To compose the 4D Gaussian field, we use B_i to define the particle field size and combine it with quantity q , which controls the number of weather particles to regulate density. The final attributes of the 4D Gaussian field, are denoted as: $W_i = \{G(A_i), q, B_i\}$.

Dynamic Simulation. Real-world weather is dynamic, with particles falling under gravity and wind. Instead of simulating these forces in detail, we approximate motion using a constant directional velocity, updating each particle's posi-

tion as follows:

$$P_j(t + \Delta t) = P_j(t) + D_j \cdot \Delta t, \quad (5)$$

where

$$P_j(t) = [x(t), y(t), z(t)]^\top, D_j = [D_x, D_y, D_z]^\top \quad (6)$$

represent the position and velocity of the j -th particle, respectively. Here, $x(t), y(t), z(t)$ denote the spatial coordinates of the particle at time t , while D_x, D_y, D_z are its velocity components along the x, y , and z axes.

To efficiently simulate a consistent falling particle effect, we recycle particles when they exit the predefined simulation field by resetting their positions. The position of the i -th particle, $P_i(t)$, is updated as follows:

$$P_i(t) = \begin{cases} p_{\max} - \delta_p, & \text{if } p_i(t) < p_{\min}, \\ p_{\min} + \delta_p, & \text{if } p_i(t) > p_{\max}, \\ P_i(t), & \text{otherwise.} \end{cases} \quad (7)$$

Here, p_{\min} and p_{\max} denotes the field’s lower and upper bounds, while δ_p represents the offset to the field bounds.

Local Field Alignment. After obtaining the reconstructed 3D scene and simulated 4D Gaussian field, we combine them to render the final weather scene. In large-scale environments, the weather field must scale appropriately to ensure a consistent weather effect. However, as scene size increases, more particles are required, leading to higher computational costs. To mitigate this, we constrain the size of the 4D Gaussian field and limit the number of particles, aligning its motion with the rendering camera. This keeps the field relatively stationary to the camera while preserving internal dynamics and intended weather effects. Given camera poses \mathbf{T}_0 at t_0 and \mathbf{T}_t at t , the relative transformation is:

$$\Delta \mathbf{T} = \mathbf{T}_t \mathbf{T}_0^{-1}. \quad (8)$$

The particle’s position, initially $P_j(t_0)$, updates as:

$$P_j(t) = \Delta \mathbf{R} P_j(t_0) + \Delta \mathbf{t}, \quad (9)$$

where $\Delta \mathbf{R}$ and $\Delta \mathbf{t}$ are the rotation and translation from $\Delta \mathbf{T}$. In this way, we can ensure consistent weather dynamics in large-scale scenes with minimal computational cost.

Experiments

Experimental Setup

Dataset. To develop the weather dataset for all-in-one adapter fine-tuning, we extract 1,237 image pairs from the BDD100K (Yu et al. 2020), MUSE (Brödermann et al. 2024), and ACDC (Sakaridis, Dai, and Van Gool 2021) datasets. Each pair contains a normal-weather image and an adverse-weather image, including snowy, rainy, and foggy conditions, along with a corresponding ground-truth semantic segmentation map. For weather editing, we selected eight scenes with multi-view image sequences from the Pandaset (Xiao et al. 2021), Waymo Open Dataset (Sun et al. 2020), nuScenes (Caesar et al. 2020), and KITTI-360 (Liao, Xie, and Geiger 2022) datasets.

Implementation details. We implement the 2D editing model using CycleGAN-Turbo (Parmar et al. 2024), using a

	ControlNet	IP2P	TurboEdit	Ours
CLIP-DS Score (\uparrow)	0.291	0.295	0.257	0.302
CLIP-S Score (\uparrow)	0.594	0.691	0.676	0.750

Table 1: Comparison with 2D-based weather editing.

	3D Stylization	ClimateNeRF	Ours
CLIP-DS Score (\uparrow)	0.225	0.267	0.283
CLIP-S Score (\uparrow)	0.842	0.702	0.751

Table 2: Comparison with 3D-based weather editing.

single RTX A6000, which takes 12 hours training and runs at 3.2 seconds per frame. The 3D scene reconstruction and 4D Gaussian field are based on OmniRe (Chen et al. 2025) and 3DGS (Kerbl et al. 2023), with both training (~1 hour) and rendering (~0.21s) performed on a single RTX 3090.

Evaluation Metrics. To evaluate edited weather background quality, we use the cosine similarity of CLIP (Radford et al. 2021) image embeddings (CLIP-S) to measure content preservation and the directional CLIP similarity (Gal et al. 2022) (CLIP-DS) to assess alignment with text instructions. To assess consistency across temporal and spatial dimensions, we calculate the warp error to measure temporal coherence between adjacent frames and employ the Bhattacharyya distance (Bhattacharyya 1943) to quantify color distribution similarity across viewpoints.

Results

Experimental Results

We present both qualitative and quantitative comparisons of our weather background editing method against 2D-based approaches, including ControlNet (Zhang, Rao, and Agrawala 2023), InstructPix2Pix (Brooks, Holynski, and Efros 2023), and TurboEdit (Deutch et al. 2024), as well as 3D-based methods, such as 3D stylization (Li et al. 2018) and ClimateNeRF (Li et al. 2023). For dynamic weather effects, we apply qualitative analysis, *i.e.*, controllable weather simulation, to demonstrate the results.

Weather Background Editing Comparison. For the comparison of 2D-based methods, as shown in Figure 6, we observe that ControlNet tends to over-stylize the image, altering its original content. InstructPix2Pix produces a reasonable effect; however, the effect is inconsistent across adjacent frames and viewpoints. While TurboEdit accurately stylizes the image, it significantly modifies the scene content (e.g., turning pedestrians into water reflections and trees in a foggy scene). In contrast, our method achieves a realistic and coherent weather transformation while preserving scene integrity, as also indicated by the highest CLIP-DS and CLIP-S scores shown in Table 1.

For 3D-based methods, as shown in Figure 7, 3D stylization produces only a barely perceptible snowy effect on the vegetation, aligning with the highest CLIP-S score (shown in Table 2), which indicates minimal change in the scene. ClimateNeRF fails to modify the overall scene tone to ac-

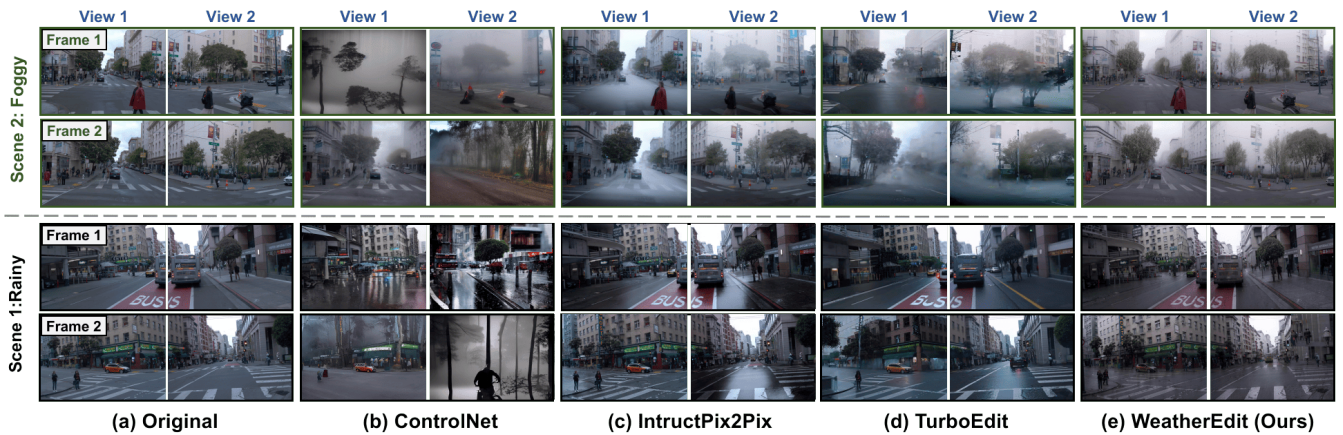


Figure 6: Comparison with 2D-based Editing Methods. We show (a) original images with edited results from (b) ControlNet, (c) InstructPix2Pix, (d) TurboEdit, and (e) Ours. Existing methods often suffer from over-stylization, content removal, and inconsistencies in spatial and temporal coherence. In contrast, our method preserves scene integrity while ensuring multi-view and multi-frame consistency, producing realistic and temporally coherent weather effects across viewpoints and frames.



Figure 7: Comparison with 3D-based weather synthesis methods in editing *snowy*. 3D Stylization (b) produces only a subtle snowy effect on the sidewalk, while (c) ClimateNeRF fails to adjust overall scene tone. In contrast, our method (d) synthesizes a realistic snowy effect while ensuring a perceptually convincing atmosphere.




Figure 8: Controllable Weather Effect and Severity in Driving Scenes. WeatherEdit can modify (a) original driving scenes by generating weather effects while preserving scene structure in (b). Additionally, it enables (c) controllable weather severity, adjusting precipitation intensity from light to heavy for realistic and flexible weather simulation.

curately reflect the intended weather conditions. In contrast, our approach not only synthesizes a realistic snowy effect but also adjusts the scene tone to create a coherent and perceptually convincing snowy atmosphere, as also indicated by the highest CLIP-DS score in Table 2.

Controllable Weather Simulation. A key advantage of our method over existing approaches is its ability to generate dynamic weather effects while precisely controlling both the type and severity of weather conditions. Additionally, we can flexibly choose whether to introduce weather particles

	ACDC		MUSE	
	w/o	w	w/o	w
HRDA	46.3	58.7 (+12.4)	42.2	57.1 (+14.9)
MIC	50.0	59.0 (+9.0)	48.8	55.4 (+6.8)

Table 3: Comparison of segmentation results with and without our simulated weather data.

and, if included, adjust their intensity with fine granularity. This allows for a range of scenarios, from a static wet road surface without active rainfall to a heavy snowstorm with dense snowfall obscuring visibility. To visualize the results, we render edited weather backgrounds and modulate weather severity. As shown in Figure 8b, backgrounds are reconstructed from 2D editing outputs, with weather type determined by the selected adapters. We then apply the 4D Gaussian field to introduce particles (Figure 8c). By adjusting particle attributes such as color, opacity, shape, rotation, and velocity, we simulate snow, rain, and fog from light to heavy. This enables flexible, physically consistent scene manipulation for diverse weather conditions.

Benefits for Downstream Tasks. We validate the benefits of our simulation as data augmentation for a downstream semantic segmentation task on adverse weather dataset such as ACDC(Sakaridis, Dai, and Van Gool 2021), MUSE(Brödermann et al. 2024). We use HRDA(Hoyer, Dai,




Figure 9: Ablation on WeatherEdit modules.



Figure 10: Effect of Temporal-View Attention on Consistency.

and Van Gool 2022) and MIC(Hoyer et al. 2023) for training. As shown in Table 3, our simulated data boosts the performance across both datasets and models, with mIoU gains up 14.9%. This shows that we can enhance model robustness under adverse weather conditions, validating the utility of WeatherEdit for downstream tasks.

Ablation Studies

We conduct ablation studies on the semantic condition input for finetuning, temporal-view attention for image editing and modular design in weather particle construction.

Semantic Conditioned Input. To evaluate the effectiveness of semantic-conditioned input, we compare editing results and FID scores with and without it. By computing the cross-attention map using “snowy” in the text prompt, we observe that semantic conditioning enhances the model’s response to sidewalks and trees, aligning weather effects more accurately with scene structures, as shown in Figure 11. Additionally, the model achieves a lower average FID value across different weather conditions compared to cases without semantic conditioning, demonstrating its effectiveness in capturing weather characteristics during training.

Temporal-View Attention. As shown in Figure 10, without TV-Attention, view 3 exhibits inconsistent editing effects, with a darker tone across two frames compared to view 2. After incorporating temporal-view attention, view 3 achieves more consistent editing, displaying a similar snowy effect between adjacent views and frames. For the qualitative analysis shown in Table 4, temporal attention reduces warp error, indicating improved temporal coherence, while spatial attention results in a lower Bhattacharyya distance, suggesting more consistent editing across viewpoints. By integrating both temporal and view attention, the full implementation achieves the lowest Bhattacharyya distance while maintaining a balanced reduction in warp error, ensuring globally consistent editing results.

Modular Design for 4D Gaussian Field. As the cross-frame results in Figure 9 show, without 2D background edit-




Figure 11: Ablation on semantic conditioned input. Left (a) shows the input image paired with a segmentation map. With semantic conditioning (b), snowfall presents more prominent feature (highlighted in red box) and aligns better with scene structures compared to (c) without it.

	Self-attn	+View	+Temp	Full
Warp-error (\downarrow)	0.043	0.042	0.039	0.041
Bhattacharyya dist (\downarrow)	0.272	0.253	0.261	0.245

Table 4: Ablation on temporal-view attention.

ing (b), the scene lacks an overall weather tone. Randomly initializing 4D Gaussians for attribute modeling (c) produces colorful blobs instead of realistic weather particles. Without dynamic simulation (d), the particles remain relatively static in the camera view, lacking the natural falling motion. Without local field alignment (e), weather particles stay fixed at the initial position of the scene. As the camera moves forward, the weather particles gradually disappear instead of persisting throughout the scene. In contrast, when all modules are included, the full implementation (f) generates a realistic weather background with natural dynamics.

Conclusion

In this work, we introduce WeatherEdit, a novel 4D weather editing framework capable of generating realistic weather effects such as fog, rain, and snow, with fine-grained control over their severity. Our approach leverages an all-in-one adapter that enables a diffusion model to transform normal-weather images into diverse weather styles. To maintain consistency of edited results across multiple frames and viewpoints, we introduce an ordered Temporal-View (TV) Attention mechanism. Additionally, by designing a 4D Gaussian field, we model the attributes and dynamics of weather particles within a 3D scene, enhancing realism and enabling fine-grained control over weather severity. WeatherEdit demonstrates its potential in generating extreme weather conditions from normal scenes, providing a foundation for future research on evaluating the resilience of real-world applications such as autonomous driving.

Acknowledgments

This work was funded by the National Natural Science Foundation of China (No. 62572166, No. 62402157) and the Fundamental Research Funds for the Central Universities (No. JZ2025HG7B0219).

References

- Almalioglu, Y.; Turan, M.; Trigoni, N.; and Markham, A. 2022. Deep learning-based robust positioning for all-weather autonomous driving. *Nature Machine Intelligence*, 4(9): 749–760.
- Bhattacharyya, A. K. 1943. On a measure of divergence between two statistical populations defined by their probability distributions. *Bulletin of the Calcutta Mathematical Society*, 35: 99–109.
- Brödermann, T.; Bruggemann, D.; Sakaridis, C.; Ta, K.; Liagouris, O.; Corkill, J.; and Van Gool, L. 2024. Muses: The multi-sensor semantic perception dataset for driving under uncertainty. In *European Conference on Computer Vision*, 21–38. Springer.
- Brooks, T.; Holynski, A.; and Efros, A. A. 2023. InstructPix2Pix: Learning to Follow Image Editing Instructions. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 18392–18402.
- Caesar, H.; Bankiti, V.; Lang, A. H.; Vora, S.; Liong, V. E.; Xu, Q.; Krishnan, A.; Pan, Y.; Baldan, G.; and Beijbom, O. 2020. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11621–11631.
- Chen, Z.; Yang, J.; Huang, J.; de Lutio, R.; Esturo, J. M.; Ivanovic, B.; Litany, O.; Gojcic, Z.; Fidler, S.; Pavone, M.; Song, L.; and Wang, Y. 2025. OmniRe: Omni Urban Scene Reconstruction. In *The Thirteenth International Conference on Learning Representations*.
- Deutch, G.; Gal, R.; Garibi, D.; Patashnik, O.; and Cohen-Or, D. 2024. TurboEdit: Text-Based Image Editing Using Few-Step Diffusion Models. *arXiv:2408.00735*.
- Feldman, B. E.; and O’Brien, J. F. 2002. Modeling the accumulation of wind-driven snow. In *ACM SIGGRAPH 2002 conference abstracts and applications*, 218–218.
- Gal, R.; Patashnik, O.; Maron, H.; Bermano, A. H.; Chechik, G.; and Cohen-Or, D. 2022. Stylegan-nada: Clip-guided domain adaptation of image generators. *ACM Transactions on Graphics (TOG)*, 41(4): 1–13.
- Haas, J. K. 2014. A history of the unity game engine.
- Hahner, M.; Dai, D.; Sakaridis, C.; Zaech, J.-N.; and Van Gool, L. 2019. Semantic understanding of foggy scenes with purely synthetic data. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, 3675–3681. IEEE.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Hoyer, L.; Dai, D.; and Van Gool, L. 2022. Hrda: Context-aware high-resolution domain-adaptive semantic segmentation. In *European conference on computer vision*, 372–391. Springer.
- Hoyer, L.; Dai, D.; Wang, H.; and Van Gool, L. 2023. MIC: Masked image consistency for context-enhanced domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11721–11732.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Hwang, S.; Jeon, S.; Ma, Y.-S.; and Byun, H. 2022. WeatherGAN: Unsupervised multi-weather image-to-image translation via single content-preserving UResNet generator. *Multimedia Tools and Applications*, 81(28): 40269–40288.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Transactions on Graphics*, 42(4).
- Li, Y.; Lin, Z.-H.; Forsyth, D.; Huang, J.-B.; and Wang, S. 2023. ClimateNeRF: Extreme Weather Synthesis in Neural Radiance Field. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.
- Li, Y.; Liu, M.-Y.; Li, X.; Yang, M.-H.; and Kautz, J. 2018. A closed-form solution to photorealistic image stylization. In *Proceedings of the European conference on computer vision (ECCV)*, 453–468.
- Liao, Y.; Xie, J.; and Geiger, A. 2022. Kitti-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d. volume 45, 3292–3310. IEEE.
- Lin, C.-H.; Gao, J.; Tang, L.; Takikawa, T.; Zeng, X.; Huang, X.; Kreis, K.; Fidler, S.; Liu, M.-Y.; and Lin, T.-Y. 2023. Magic3d: High-resolution text-to-3d content creation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 300–309.
- Liu, X.; Zhan, X.; Tang, J.; Shan, Y.; Zeng, G.; Lin, D.; Liu, X.; and Liu, Z. 2024. Humangaussian: Text-driven 3d human generation with gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6646–6657.
- Meng, C.; He, Y.; Song, Y.; Song, J.; Wu, J.; Zhu, J.-Y.; and Ermon, S. 2021. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2020. NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis. In *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I*, 405–421. Berlin, Heidelberg: Springer-Verlag. ISBN 978-3-030-58451-1.
- Mou, C.; Wang, X.; Xie, L.; Wu, Y.; Zhang, J.; Qi, Z.; and Shan, Y. 2024. T2I-Adapter: learning adapters to dig out more controllable ability for text-to-image diffusion models. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence, AAAI’24/IAAI’24/EAAI’24*. AAAI Press. ISBN 978-1-57735-887-9.
- Parmar, G.; Park, T.; Narasimhan, S.; and Zhu, J.-Y. 2024. One-step image translation with text-to-image models. *arXiv preprint arXiv:2403.12036*.
- Qian, C.; Guo, Y.; Li, W.; and Markkula, G. 2025a. Weath-erGS: 3D Scene Reconstruction in Adverse Weather Condi-

- tions via Gaussian Splatting. In *International Conference on Robotics and Automation(ICRA)*.
- Qian, C.; Guo, Y.; Mo, Y.; and Li, W. 2025b. WeatherDG: LLM-assisted procedural weather generation for domain-generalized semantic segmentation. *IEEE Robotics and Automation Letters*.
- Qian, C.; Rezaei, M.; Anwar, S.; Li, W.; Hussain, T.; Azarmi, M.; and Wang, W. 2025c. AllWeather-Net: Unified Image Enhancement for Autonomous Driving Under Adverse Weather and Low-Light Conditions. In *International Conference on Pattern Recognition*, 151–166. Springer.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2021. High-Resolution Image Synthesis with Latent Diffusion Models. arXiv:2112.10752.
- Rothmeier, T.; and Huber, W. 2021. Let it Snow: On the Synthesis of Adverse Weather Image Data. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, 3300–3306.
- Sakaridis, C.; Dai, D.; and Van Gool, L. 2021. ACDC: The adverse conditions dataset with correspondences for semantic driving scene understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10765–10775.
- Schmidt, V.; Luccioni, A.; Muckavilli, S. K.; Balasooriya, N.; Sankaran, K.; Chayes, J.; and Bengio, Y. 2019. Visualizing the consequences of climate change using cycle-consistent adversarial networks. *arXiv preprint arXiv:1905.03709*.
- Schmidt, V.; Luccioni, A.; Teng, M.; Zhang, T.; Reynaud, A.; Raghupathi, S.; Cosne, G.; Juraver, A.; Vardanyan, V.; and Hernández-García, A. 2022. ClimateGAN: Raising Climate Change Awareness by Generating Images of Floods. In *International Conference on Learning Representations*.
- Stomakhin, A.; Schroeder, C.; Chai, L.; Teran, J.; and Selle, A. 2013. A material point method for snow simulation. *ACM Transactions on Graphics (TOG)*, 32(4): 1–10.
- Sun, P.; Kretschmar, H.; Dotiwalla, X.; Chouard, A.; Patnaik, V.; Tsui, P.; Guo, J.; Zhou, Y.; Chai, Y.; Caine, B.; Vasudevan, V.; Han, W.; Ngiam, J.; Zhao, H.; Timofeev, A.; Ettinger, S.; Krivokon, M.; Gao, A.; Joshi, A.; Zhang, Y.; Shlens, J.; Chen, Z.; and Anguelov, D. 2020. Scalability in Perception for Autonomous Driving: Waymo Open Dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Valanarasu, J. M. J.; Yasarla, R.; and Patel, V. M. 2022. TransWeather: Transformer-Based Restoration of Images Degraded by Adverse Weather Conditions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2353–2363.
- Xiao, P.; Shao, Z.; Hao, S.; Zhang, Z.; Chai, X.; Jiao, J.; Li, Z.; Wu, J.; Sun, K.; Jiang, K.; et al. 2021. Pandaset: Advanced sensor suite dataset for autonomous driving. In *2021 IEEE international intelligent transportation systems conference (ITSC)*, 3095–3101. IEEE.
- Yan, Z.; Li, C.; and Lee, G. H. 2023. Nerf-ds: Neural radiance fields for dynamic specular objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8285–8295.
- Yang, Y.; Wu, H.; Aviles-Rivero, A. I.; Zhang, Y.; Qin, J.; and Zhu, L. 2024. Genuine Knowledge from Practice: Diffusion Test-Time Adaptation for Video Adverse Weather Removal. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 25606–25616.
- Yu, F.; Chen, H.; Wang, X.; Xian, W.; Chen, Y.; Liu, F.; Madhavan, V.; and Darrell, T. 2020. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2636–2645.
- Zhang, H.; Ba, Y.; Yang, E.; Mehra, V.; Gella, B.; Suzuki, A.; Pfahnl, A.; Chandrappa, C. C.; Wong, A.; and Kadambi, A. 2023. WeatherStream: Light Transport Automation of Single Image Deweathering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 13499–13509.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding Conditional Control to Text-to-Image Diffusion Models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 3836–3847.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, 2223–2232.
- Zsolnai-Fehér, K. 2022. The flow from simulation to reality. *nature physics*, 18(11): 1260–1261.