Restoring Vision in Adverse Weather Conditions with Patch-Based Denoising Diffusion Models

Ozan Özdenizci [1,2] and Robert Legenstein [1]

[1] Institute of Theoretical Computer Science, Graz University of Technology, Graz, Austria
[2] TU Graz - SAL Dependable Embedded Systems Lab, Silicon Austria Labs, Graz, Austria

IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) 2023

Restoring Vision in Adverse Weather Conditions with Patch-Based Denoising Diffusion Models

Ozan Özdenizci and Robert Legenstein

Abstract—Image restoration under adverse weather conditions has been of significant interest for various computer vision applications. Recent successful methods rely on the current progress in deep neural network architectural designs (e.g., with vision transformers). Motivated by the recent progress achieved with state-of-the-art conditional generative models, we present a novel patch-based image restoration algorithm based on denoising diffusion probabilistic models. Our patch-based diffusion modeling approach enables size-agnostic image restoration by using a guided denoising process with smoothed noise estimates across overlapping patches during inference. We empirically evaluate our model on benchmark datasets for image desnowing, combined deraining and dehazing, and raindrop removal. We demonstrate our approach to achieve state-of-the-art performances on both weather-specific and multi-weather image restoration, and qualitatively show strong generalization to real-world test images.

Index Terms—denoising diffusion models, patch-based image restoration, deraining, desnowing, dehazing, raindrop removal.

Introduction & Motivation

- **Problem:** Restoration of adverse weather related degradations from images.
- **Approach:** Generative DNNs trained on synthetic clean-distorted image pairs.
    - We develop a novel approach based on *denoising diffusion models*.

Introduction & Motivation

- **Problem:** Restoration of adverse weather related degradations from images.
- **Approach:** Generative DNNs trained on synthetic clean-distorted image pairs.
    - We develop a novel approach based on *denoising diffusion models*.

Our contributions are summarized as follows:

- We present a novel patch-based diffusive image restoration algorithm for arbitrary sized image processing with denoising diffusion models.
- We empirically demonstrate our approach to achieve state-of-the-art performance on both weather-specific and multi-weather restoration tasks.
- We qualitatively present strong generalization from synthetic to real-world multi-weather restoration with our generative modeling perspective.

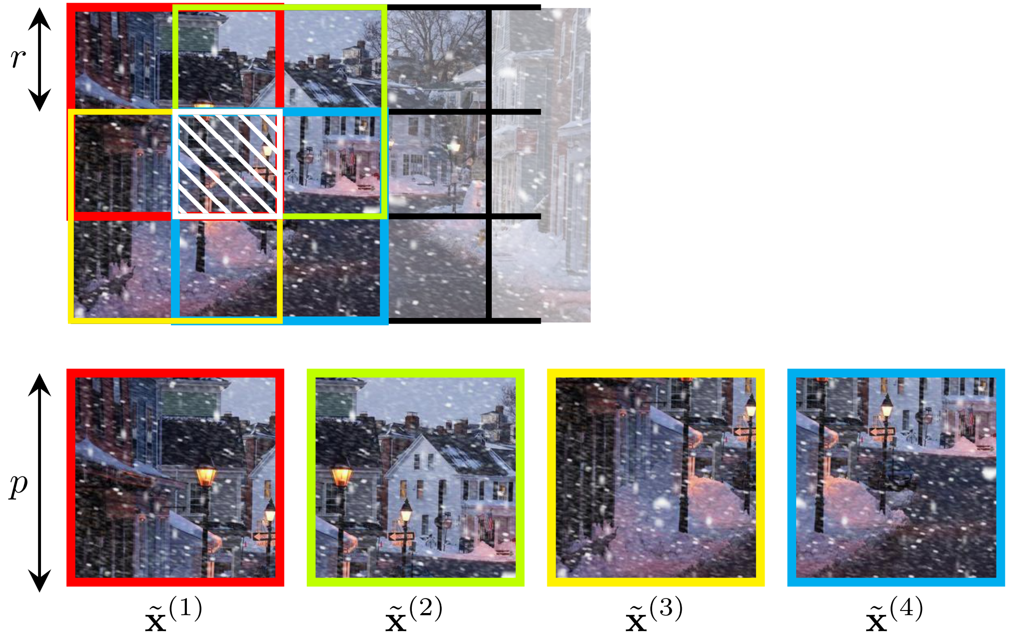

Introduction & Motivation

- **Problem:** Restoration of adverse weather related degradations from images.
- **Approach:** Generative DNNs trained on synthetic clean-distorted image pairs.
  - We develop a novel approach based on *denoising diffusion models*.

Our contributions are summarized as follows:

- We present a novel patch-based diffusive image restoration algorithm for arbitrary sized image processing with denoising diffusion models.
- We empirically demonstrate our approach to achieve state-of-the-art performance on both weather-specific and multi-weather restoration tasks.
- We qualitatively present strong generalization from synthetic to real-world multi-weather restoration with our generative modeling perspective.

Introduction & Motivation

- **Problem:** Restoration of adverse weather related degradations from images.
- **Approach:** Generative DNNs trained on synthetic clean-distorted image pairs.
    - We develop a novel approach based on *denoising diffusion models*.

Our contributions are summarized as follows:

- We present a novel patch-based diffusive image restoration algorithm for arbitrary sized image processing with denoising diffusion models.
- We empirically demonstrate our approach to achieve state-of-the-art performance on both weather-specific and multi-weather restoration tasks.
- We qualitatively present strong generalization from synthetic to real-world multi-weather restoration with our generative modeling perspective.

Background: Denoising Diffusion Probabilistic Models

Deep Unsupervised Learning using Nonequilibrium Thermodynamics

Jascha Sohl-Dickstein JASCHA@STANFORD.EDU
Stanford University
Eric A. Weiss EWEISS@BERKELEY.EDU
University of California, Berkeley
Niru Maheswaranathan NIRUM@STANFORD.EDU
Stanford University
Surya Ganguli SGANGULI@STANFORD.EDU
Stanford University

International Conference on Machine Learning (ICML) 2015

Denoising Diffusion Probabilistic Models

Jonathan Ho
UC Berkeley
jonathanho@berkeley.edu

Ajay Jain
UC Berkeley
ajayj@berkeley.edu

Pieter Abbeel
UC Berkeley
pabbeel@cs.berkeley.edu

Advances in Neural Information Processing Systems (NeurIPS) 2020

Diffusion Models Beat GANs on Image Synthesis

Prafulla Dhariwal*
OpenAI
prafulla@openai.com

Alex Nichol*
OpenAI
alex@openai.com

Advances in Neural Information Processing Systems (NeurIPS) 2021

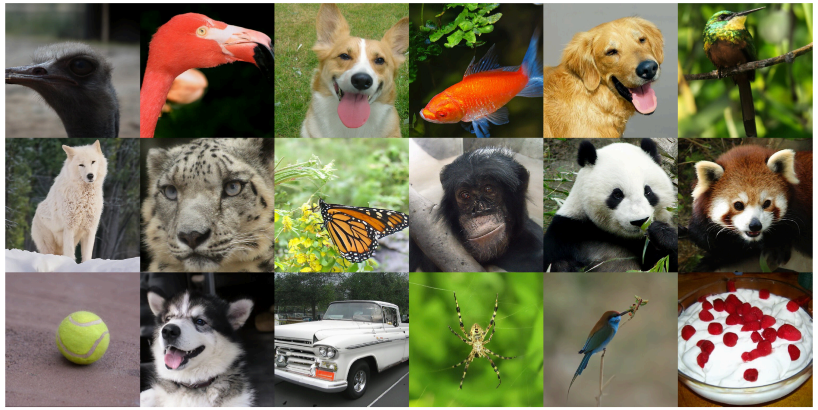

Figure 1: Selected samples from our best ImageNet 512×512 model (FID 3.85)

Background: Denoising Diffusion Probabilistic Models

<u>Summary:</u> Training a DNN that can iteratively denoise an image by reversing a diffusion process that destroys the data structure by adding noise.

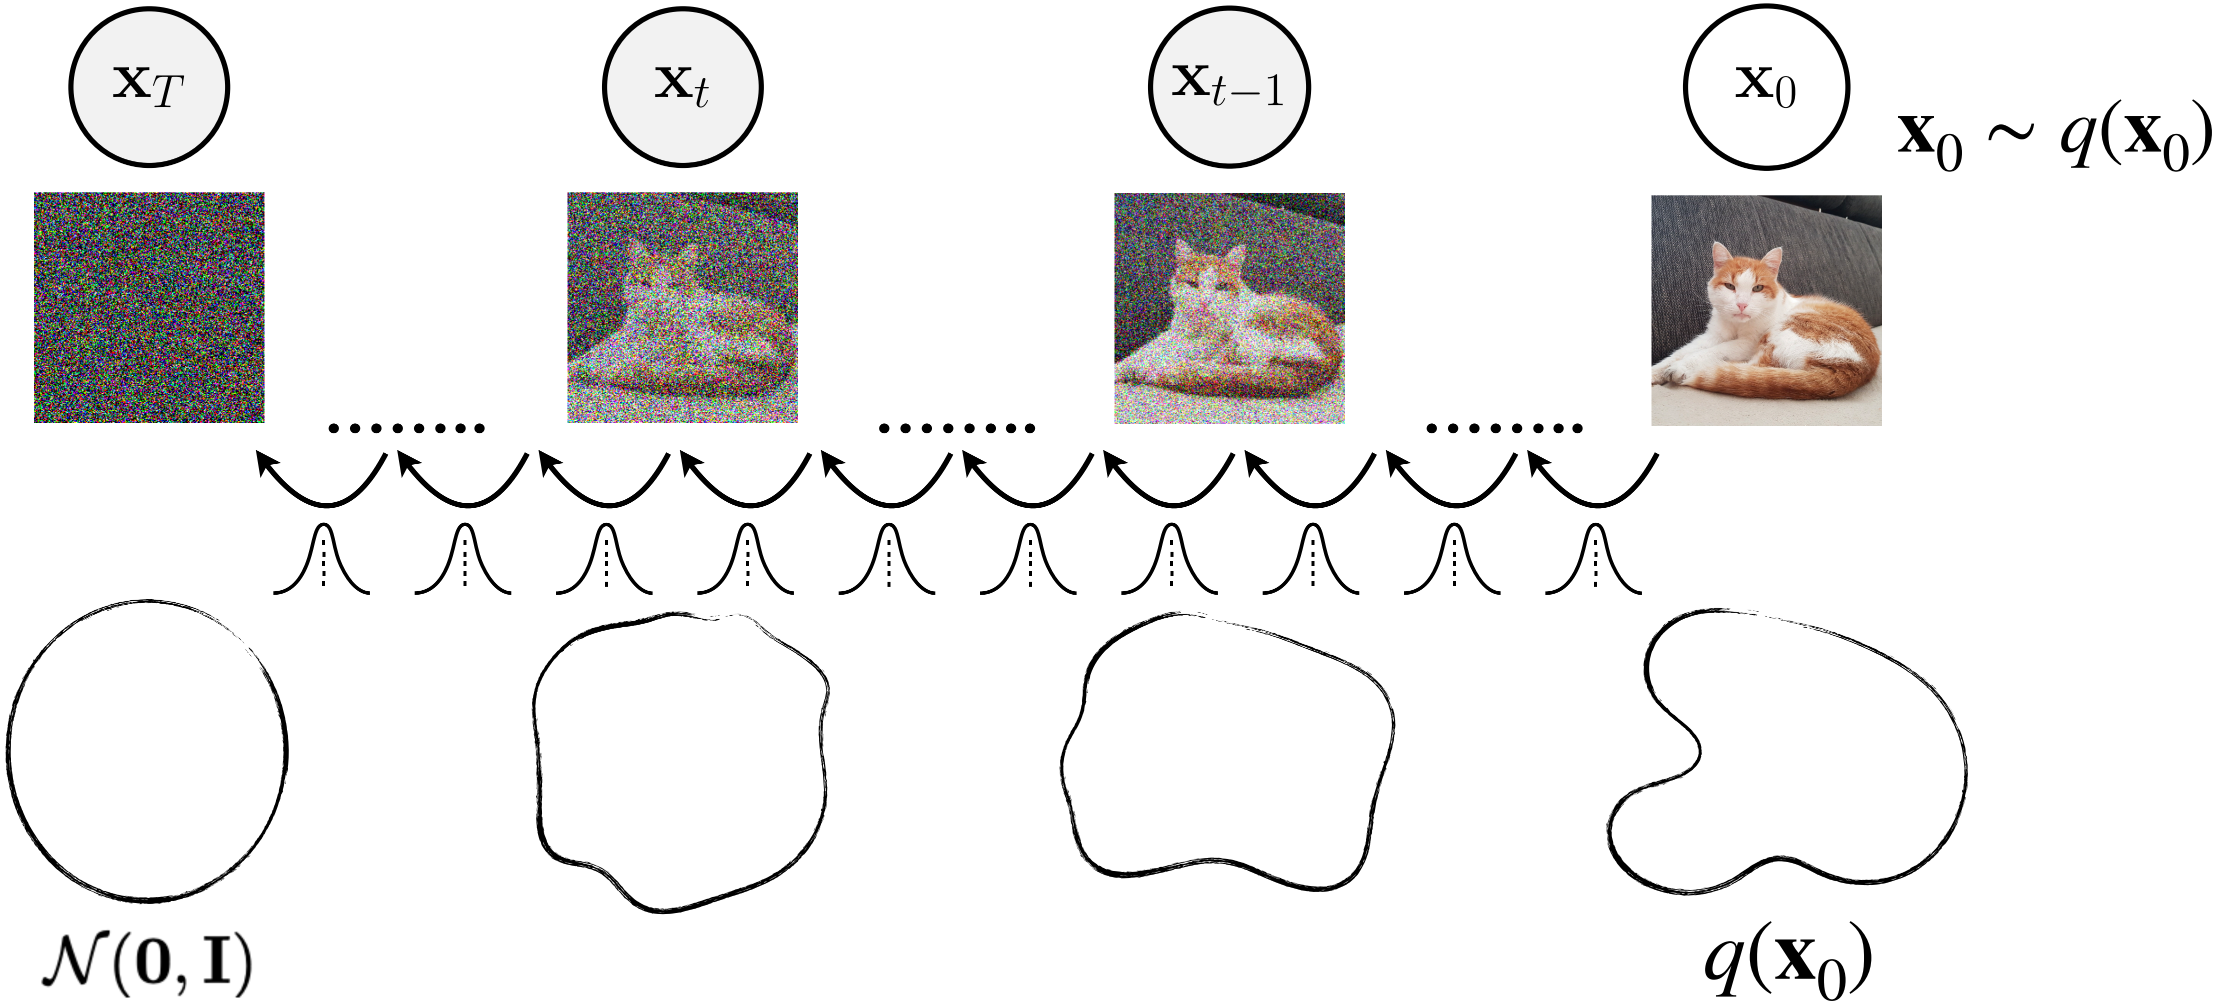

Background: Denoising Diffusion Probabilistic Models

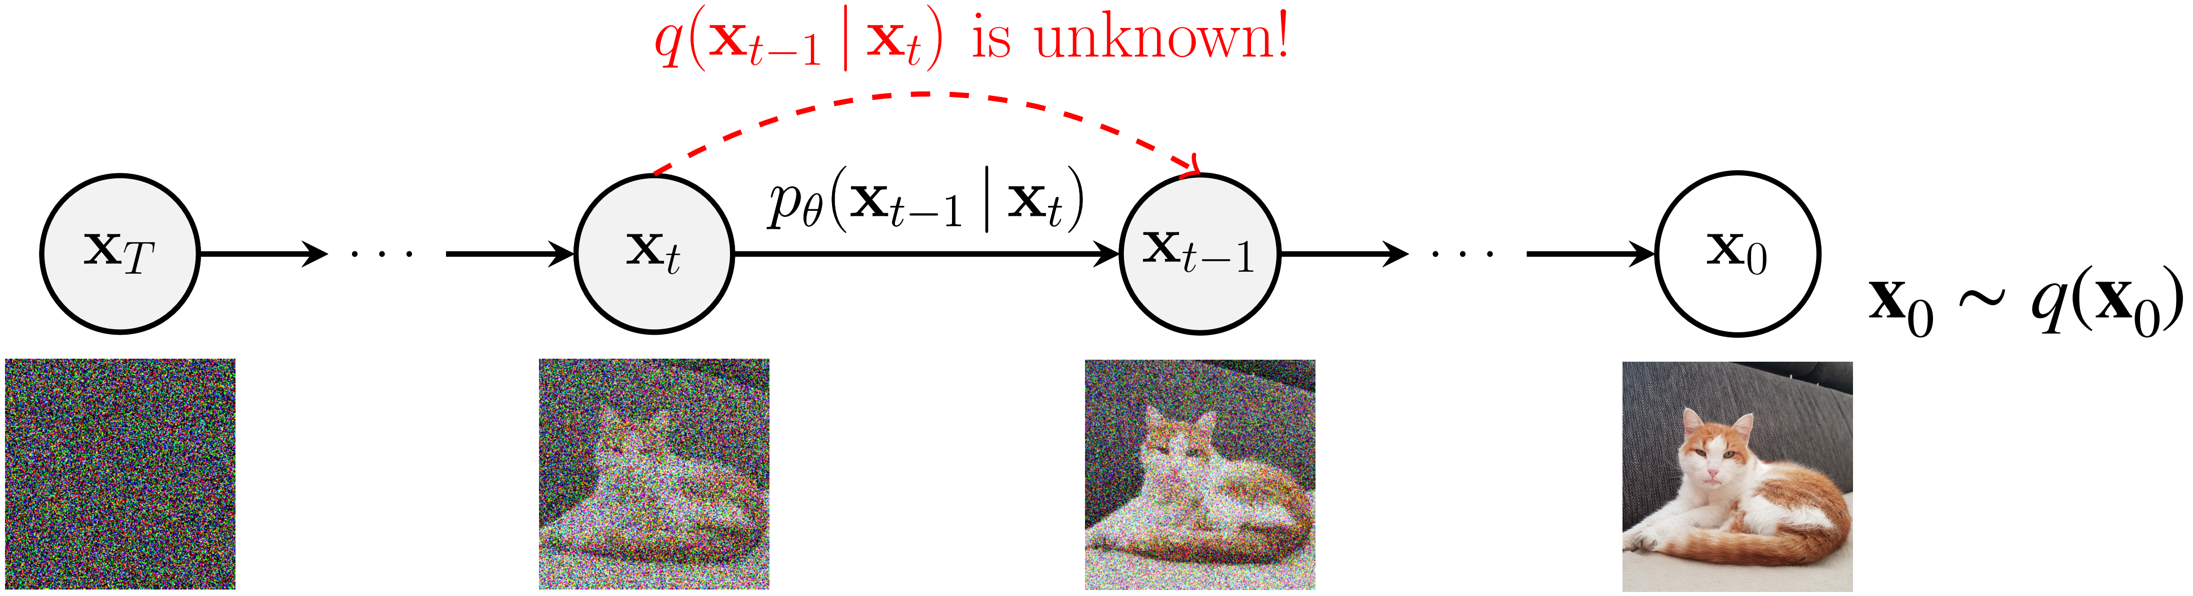

Fixed Forward (Diffusion) Process

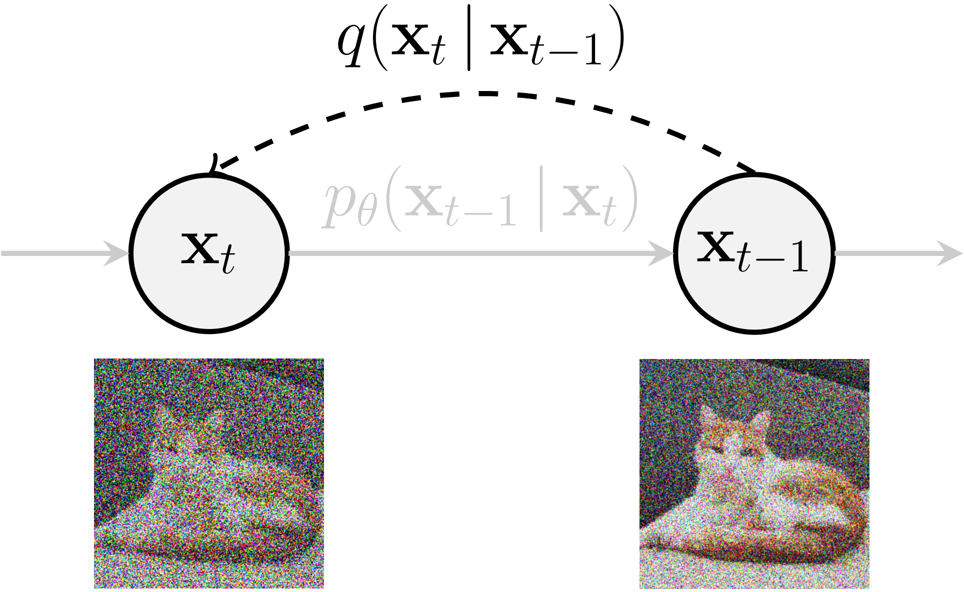

- The ***forward*** process (i.e., ***diffusion*** process) gradually adds Gaussian noise according to a known variance schedule $\beta_1 < \beta_2 < \ldots < \beta_T$.

$$q(\mathbf{x}_t \,|\, \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1-\beta_t}\,\mathbf{x}_{t-1}, \beta_t\,\mathbf{I})$$

$$q(\mathbf{x}_{1:T} \,|\, \mathbf{x}_0) = \prod_{t=1}^{T} q(\mathbf{x}_t \,|\, \mathbf{x}_{t-1}) \longrightarrow \text{joint distribution}$$

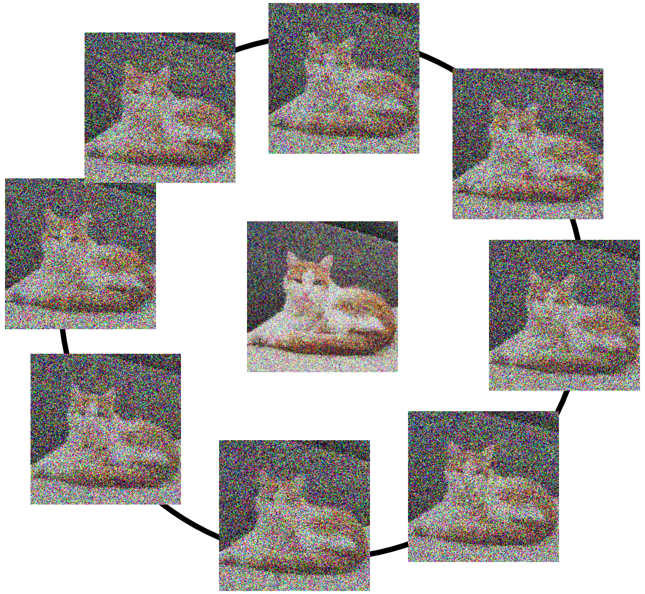

Fixed Forward (Diffusion) Process

forward process

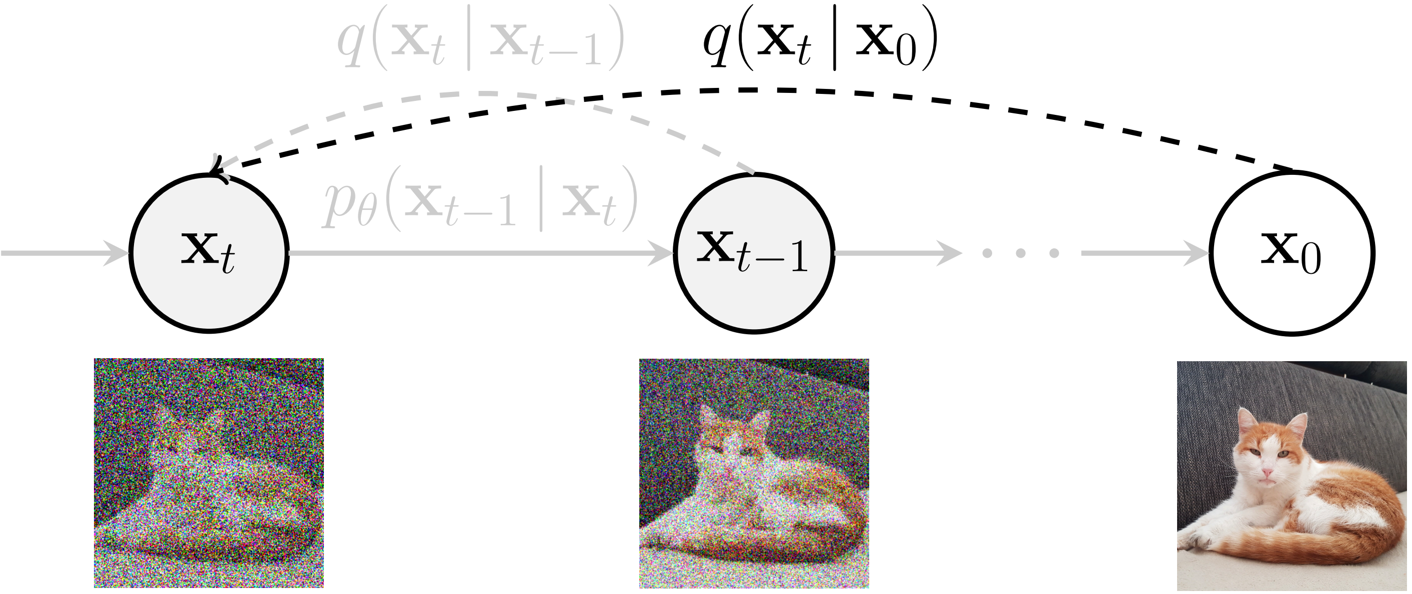

- We can also directly jump to any time-step using:

$$\alpha_t = 1 - \beta_t \quad \text{and} \quad \bar{\alpha}_t = \prod_{s=1}^{t} \alpha_s$$

$$\boxed{q(\mathbf{x}_t \mid \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t}\,\mathbf{x}_0, (1 - \bar{\alpha}_t)\,\mathbf{I})}$$

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t}\,\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\boldsymbol{\epsilon} \quad \text{where} \quad \boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

the noise schedule is designed such that:

$$\bar{\alpha}_T \to 0$$

$$\bar{\alpha}_1 > \cdots > \bar{\alpha}_T$$

Generative Reverse (Denoising) Process

reverse process

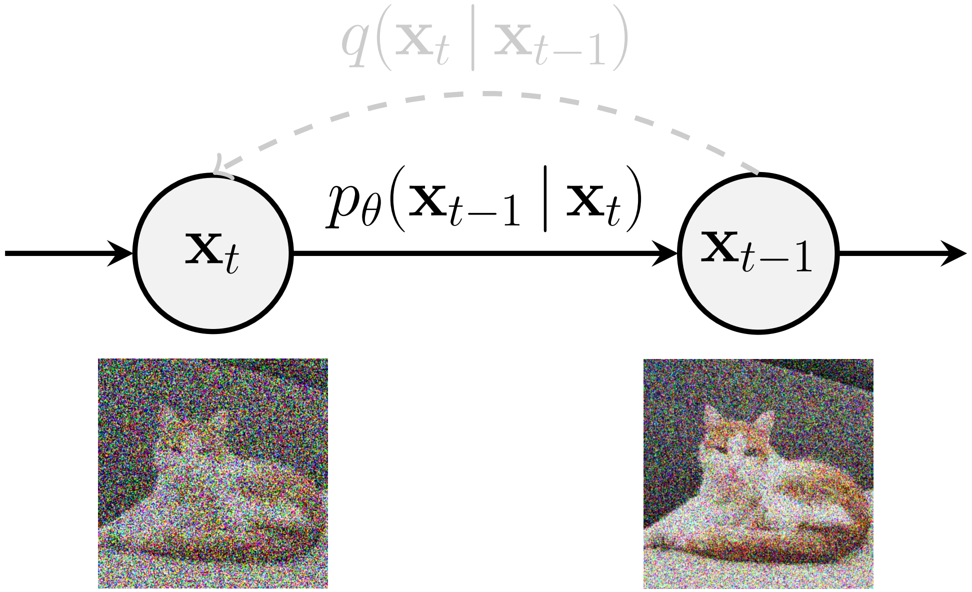

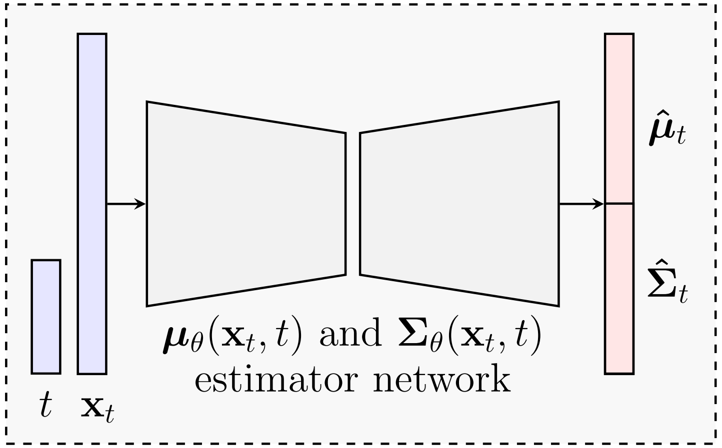

- The ***reverse*** process is the joint distribution, with learned Gaussian transitions starting from noise.

$$p_\theta(\mathbf{x}_{0:T}) = \underbrace{p_\theta(\mathbf{x}_T)}_{\text{known}} \prod_{t=1}^{T} \underline{p_\theta(\mathbf{x}_{t-1} \mid \mathbf{x}_t)}$$

$$\underline{p(\mathbf{x}_T)} = \mathcal{N}(\mathbf{x}_T; \mathbf{0}, \mathbf{I}) \qquad \underline{p_\theta(\mathbf{x}_{t-1} \mid \mathbf{x}_t)} = \mathcal{N}(\mathbf{x}_{t-1}; \mu_\theta(\mathbf{x}_t, t), \Sigma_\theta(\mathbf{x}_t, t))$$

Summary: Denoising Diffusion Probabilistic Models

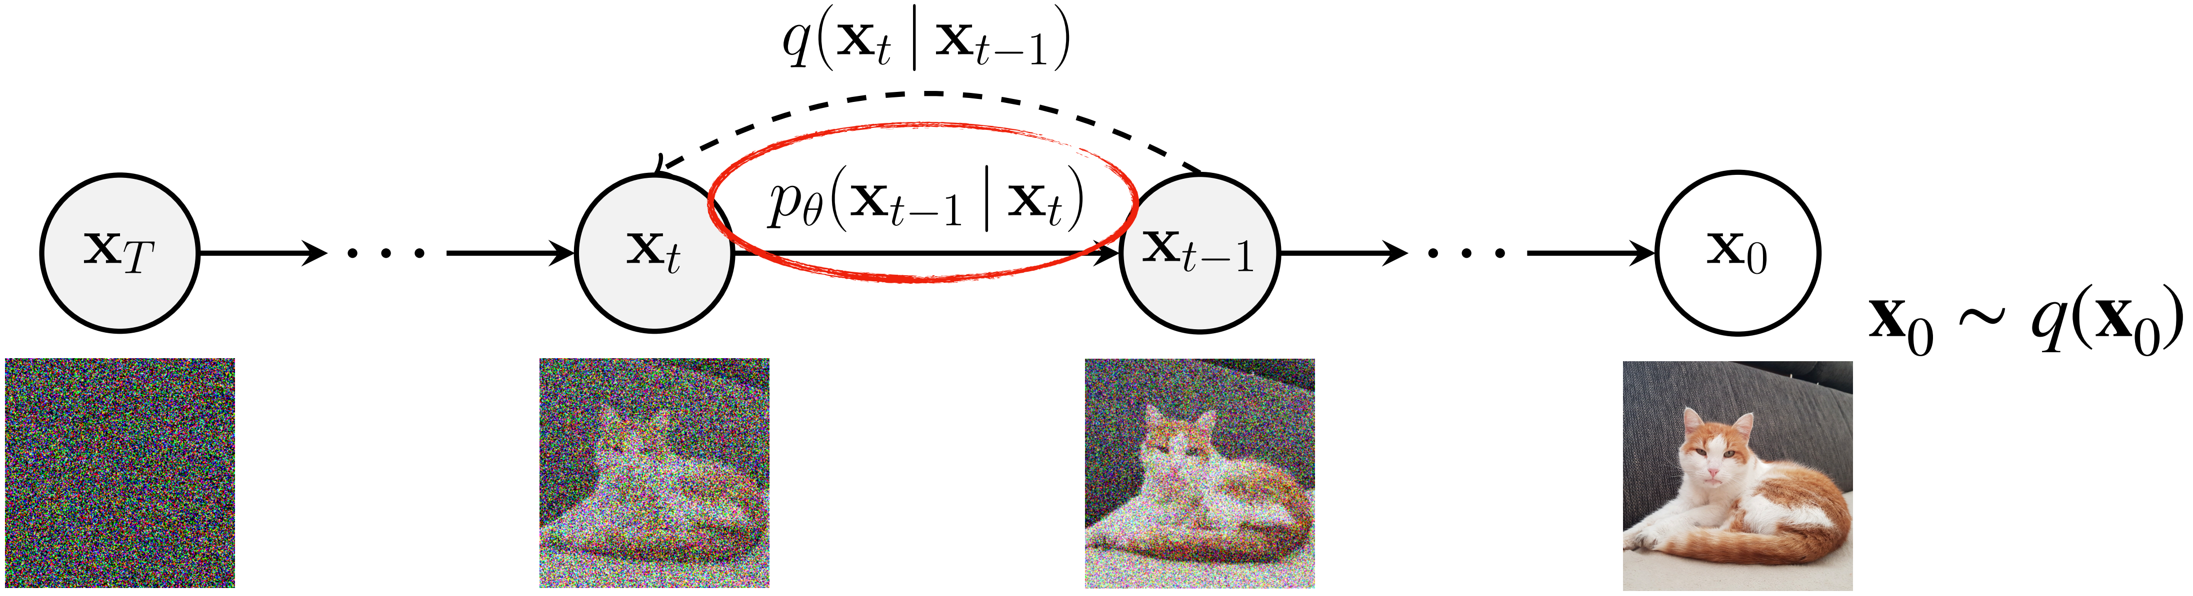

forward process

$$q(\mathbf{x}_t \mid \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1-\beta_t}\,\mathbf{x}_{t-1}, \beta_t \mathbf{I})$$

$$q(\mathbf{x}_t \mid \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t}\,\mathbf{x}_0, (1-\bar{\alpha}_t)\mathbf{I})$$

reverse process

$$p_\theta(\mathbf{x}_{t-1} \mid \mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \mu_\theta(\mathbf{x}_t, t), \Sigma_\theta(\mathbf{x}_t, t))$$

Summary: Denoising Diffusion Probabilistic Models

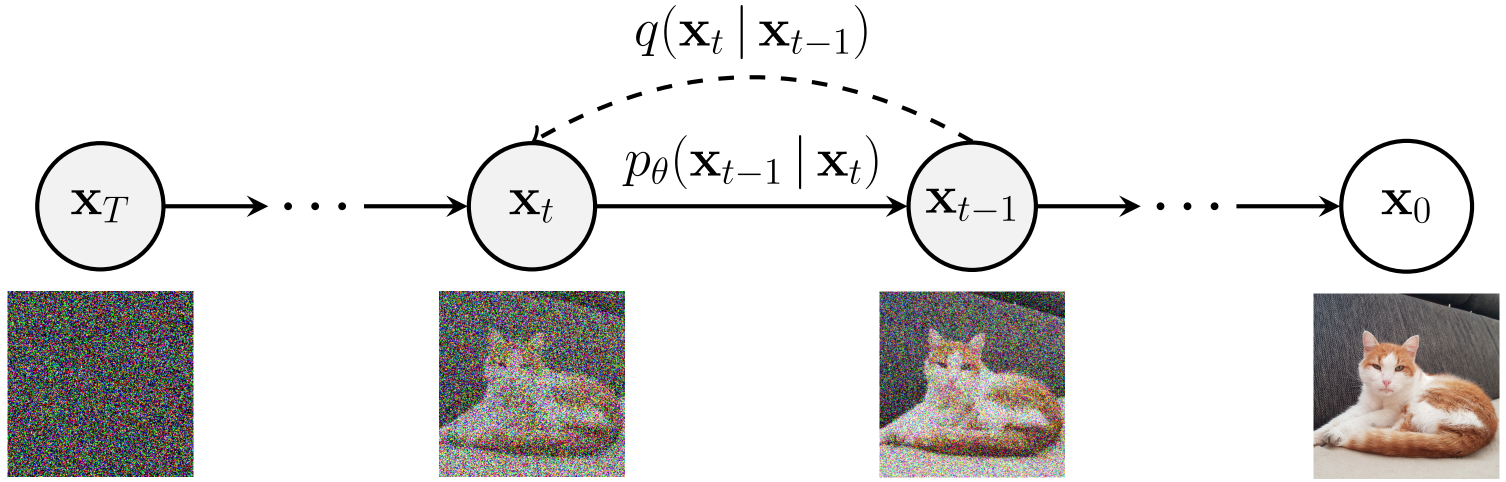

Algorithm 1 Training

1: **repeat**
2: $\mathbf{x}_0 \sim q(\mathbf{x}_0)$
3: $t \sim \mathrm{Uniform}(\{1, \ldots, T\})$
4: $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
5: Take gradient descent step on
$$\nabla_\theta \left\| \boldsymbol{\epsilon} - \boldsymbol{\epsilon}_\theta(\sqrt{\bar\alpha_t}\mathbf{x}_0 + \sqrt{1-\bar\alpha_t}\boldsymbol{\epsilon}, t) \right\|^2$$
6: **until** converged

Optimizing model parameters that predicts the added noise between ***t*** and ***t-1***

Algorithm 2 Sampling

1: $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
2: **for** $t = T, \ldots, 1$ **do**
3: $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $t > 1$, else $\mathbf{z} = \mathbf{0}$
4: $\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}}\left(\mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\bar\alpha_t}}\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)\right) + \sigma_t \mathbf{z}$
5: **end for**
6: **return** $\mathbf{x}_0$

Sequentially remove the estimated "added noise" starting from an image sampled from a Gaussian noise distribution.

Patch-based Diffusive Image Restoration: Training

- **Our approach:** Training an *image-conditional* & *patch-based* diffusion model to enable size agnostic image restoration.

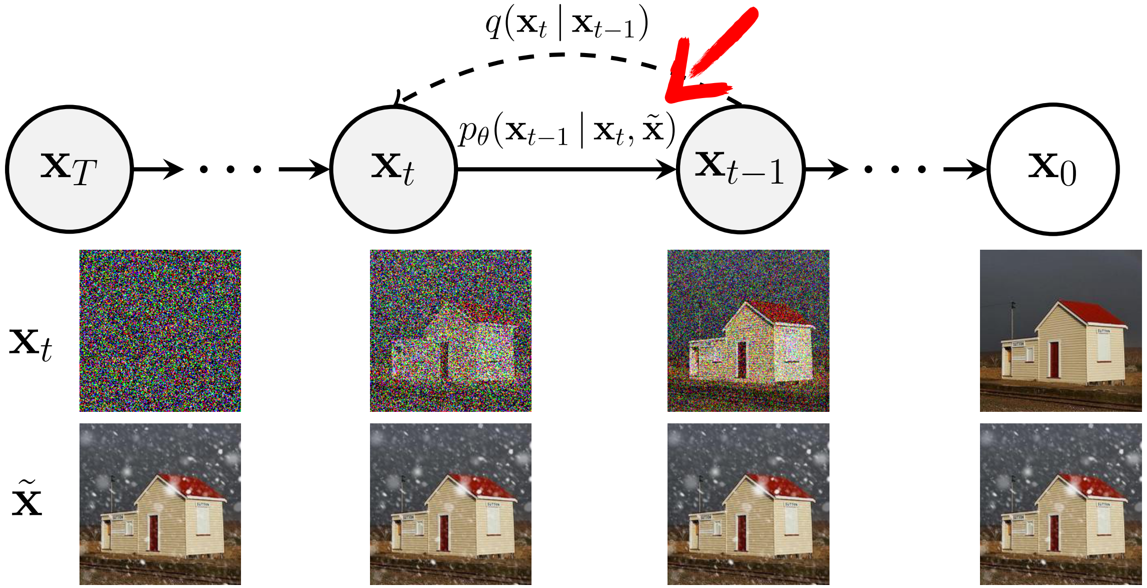

Patch-based Diffusive Image Restoration: Training

- **Our approach:** Training an <u>*image-conditional*</u> & <u>*patch-based*</u> diffusion model to enable size agnostic image restoration.

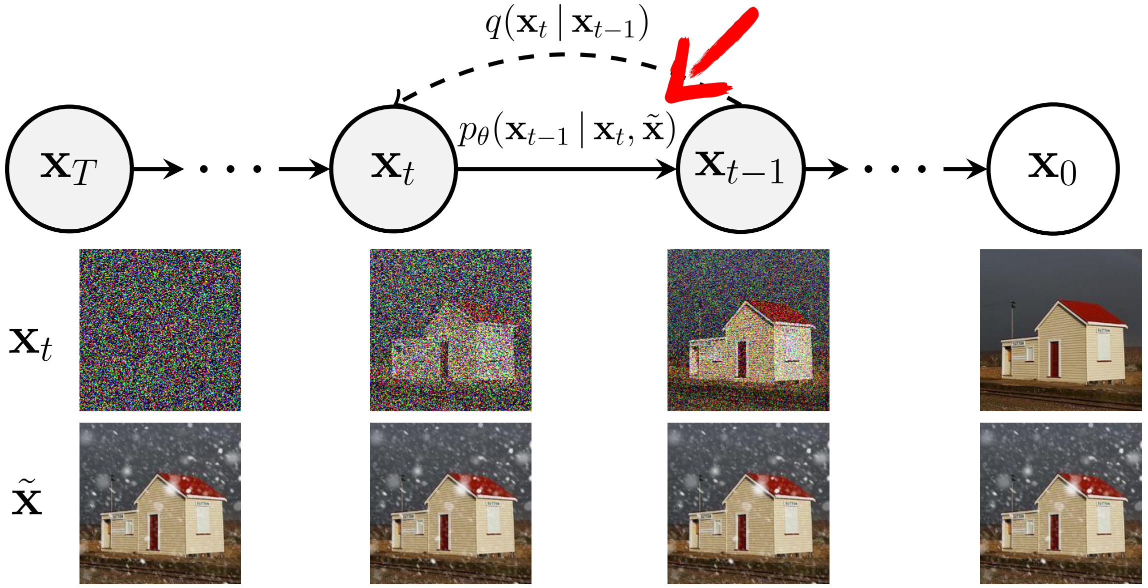

- Comes down to a simple learning algorithm that **needs low GPU memory to train and evaluate** such a model.

Algorithm 1 Diffusive weather restoration model training

Input: Clean and weather-degraded image pairs $(\mathbf{X}_0, \tilde{\mathbf{X}})$

repeat

Randomly sample a binary patch mask $\mathbf{P}_i$

$\mathbf{x}_0^{(i)} = \text{Crop}(\boldsymbol{P}_i \circ \mathbf{X}_0)$ and $\tilde{\mathbf{x}}^{(i)} = \text{Crop}(\mathbf{P}_i \circ \tilde{\mathbf{X}})$

$t \sim \text{Uniform}\{1, \ldots, T\}$

$\boldsymbol{\epsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$

Perform a single gradient descent step for

$$\nabla_\theta \| \boldsymbol{\epsilon}_t - \boldsymbol{\epsilon}_\theta(\sqrt{\bar{\alpha}_t}\,\mathbf{x}_0^{(i)} + \sqrt{1-\bar{\alpha}_t}\boldsymbol{\epsilon}_t\,, \tilde{\mathbf{x}}^{(i)}, t)\|^2$$

until converged

return θ

Patch-based Diffusive Image Restoration: Inference

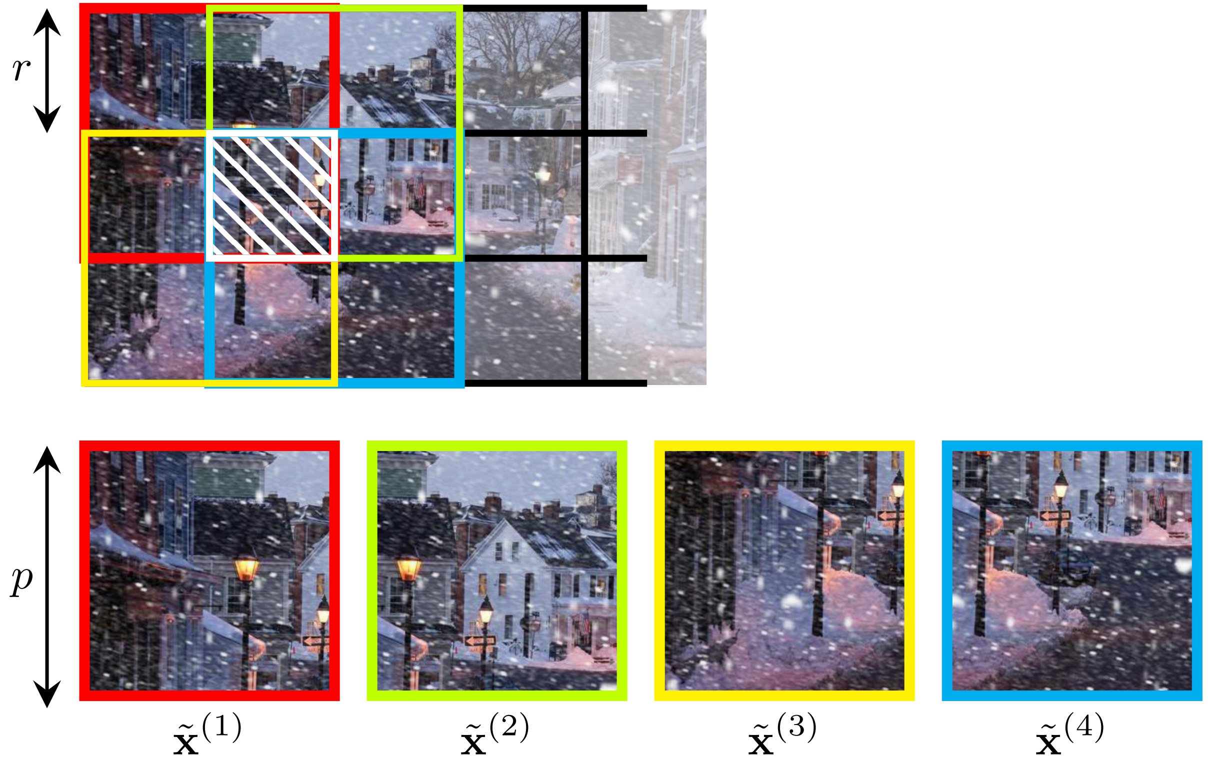

How to merge restored patches into a whole image?

Patch-based Diffusive Image Restoration: Inference

$\tilde{\mathbf{x}}^{(1)}$ $\tilde{\mathbf{x}}^{(2)}$ $\tilde{\mathbf{x}}^{(3)}$ $\tilde{\mathbf{x}}^{(4)}$

How to merge restored patches into a whole image?

Algorithm 2 Patch-based diffusive image restoration

Input: Weather-degraded image $\tilde{\mathbf{X}}$, conditional diffusion model $\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, \tilde{\mathbf{x}}, t)$, number of implicit sampling steps S, dictionary of D overlapping patch locations.

$\mathbf{X}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
for $i = S, \ldots, 1$ **do**
  $t = (i-1) \cdot T/S + 1$
  $t_{\text{next}} = (i-2) \cdot T/S + 1$ **if** $i > 1$ **else** 0
  $\hat{\boldsymbol{\Omega}}_t = \mathbf{0}$ and $\mathbf{M} = \mathbf{0}$
  for $d = 1, \ldots, D$ **do**
    $\mathbf{x}_t^{(d)} = \text{Crop}(\mathbf{P}_d \circ \mathbf{X}_t)$ and $\tilde{\mathbf{x}}^{(d)} = \text{Crop}(\mathbf{P}_d \circ \tilde{\mathbf{X}})$
    $\hat{\boldsymbol{\Omega}}_t = \hat{\boldsymbol{\Omega}}_t + \mathbf{P}_d \cdot \boldsymbol{\epsilon}_\theta(\mathbf{x}_t^{(d)}, \tilde{\mathbf{x}}^{(d)}, t)$
    $\mathbf{M} = \mathbf{M} + \mathbf{P}_d$
  end for
  $\hat{\boldsymbol{\Omega}}_t = \hat{\boldsymbol{\Omega}}_t \oslash \mathbf{M}$ // $\oslash$: element-wise division
  $\mathbf{X}_t \leftarrow \sqrt{\bar{\alpha}_{t_{\text{next}}}} \left(\frac{\mathbf{X}_t - \sqrt{1-\bar{\alpha}_t} \cdot \hat{\boldsymbol{\Omega}}_t}{\sqrt{\bar{\alpha}_t}} \right) + \sqrt{1-\bar{\alpha}_{t_{\text{next}}}} \cdot \hat{\boldsymbol{\Omega}}_t$
end for
return $\mathbf{X}_t$

Patch-based Diffusive Image Restoration: Inference

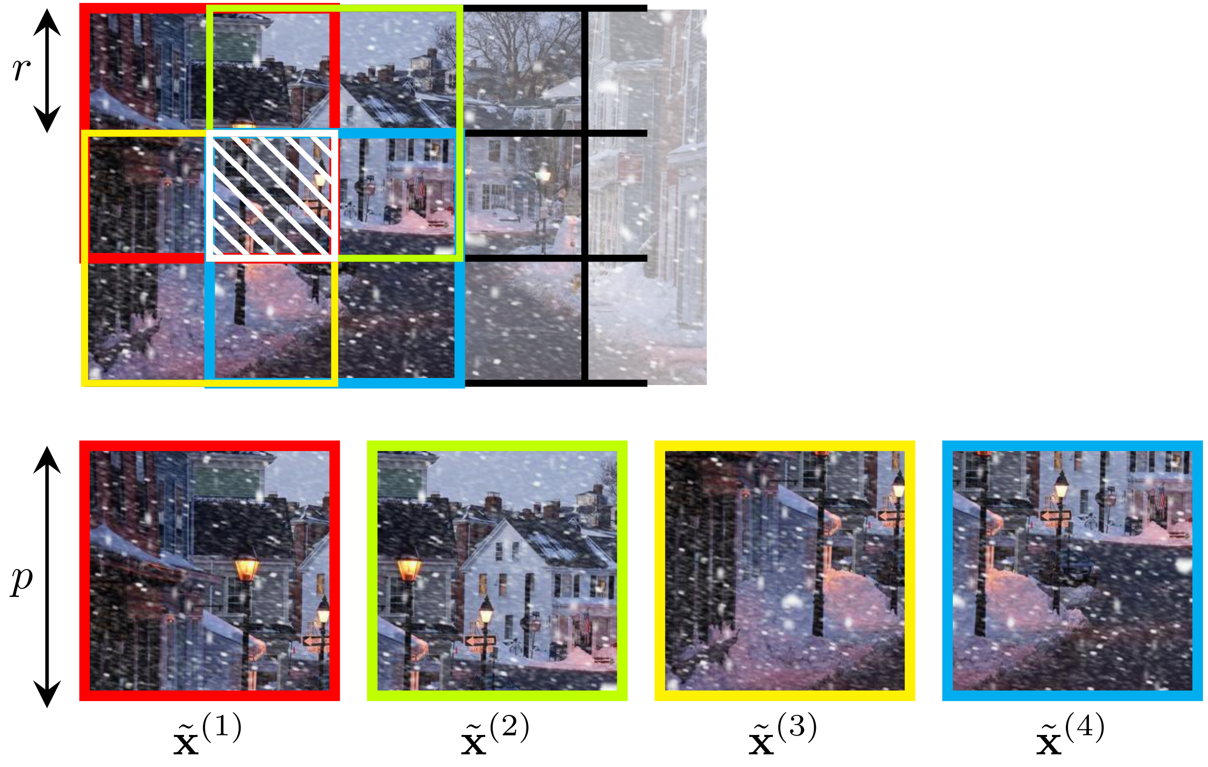

How to merge restored patches into a whole image?

At each sampling time ***t*** = ***T***, … ,***1***

1. use the "noise estimator" network for all overlapping patches to estimate the added noise at time ***t***

Patch-based Diffusive Image Restoration: Inference

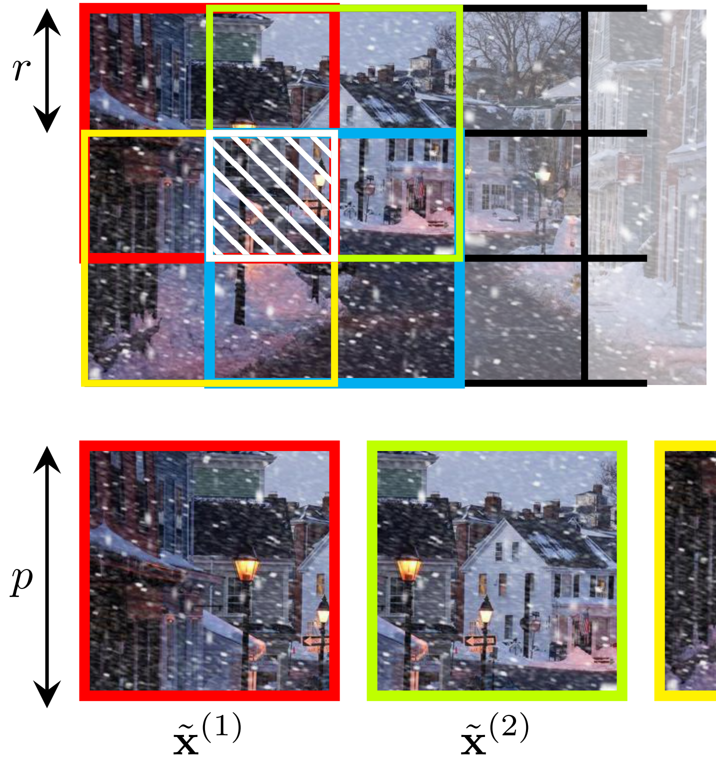

Mean estimated noise based sampling updates for the overlapping pixels:

$\frac{1}{4}\sum_{d=1}^{4}\boldsymbol{\epsilon}_\theta(\mathbf{x}_t^{(d)}, \tilde{\mathbf{x}}^{(d)}, t)$

How to merge restored patches into a whole image?

At each sampling time ***t*** = ***T***, … ,***1***

1. use the "noise estimator" network for all overlapping patches to estimate the added noise at time ***t***

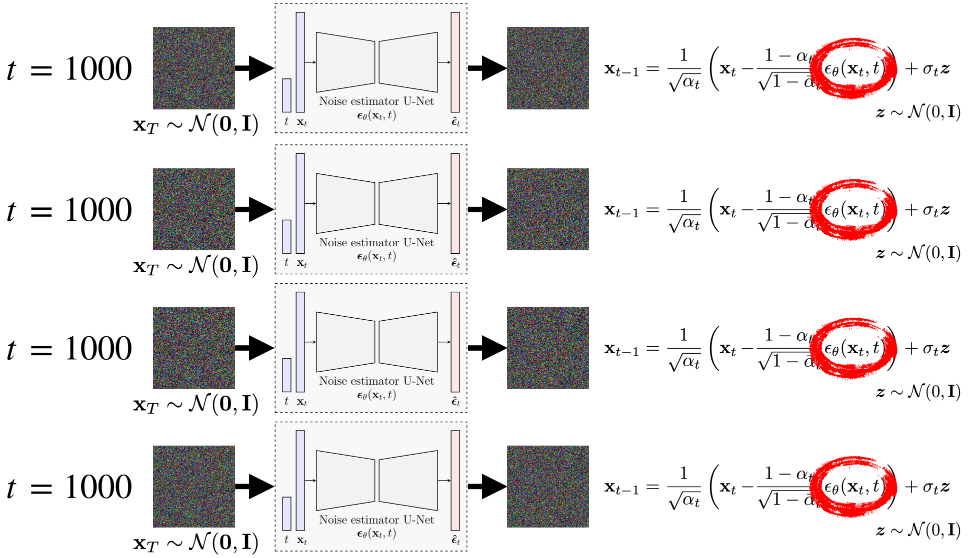

Patch-based Diffusive Image Restoration: Inference

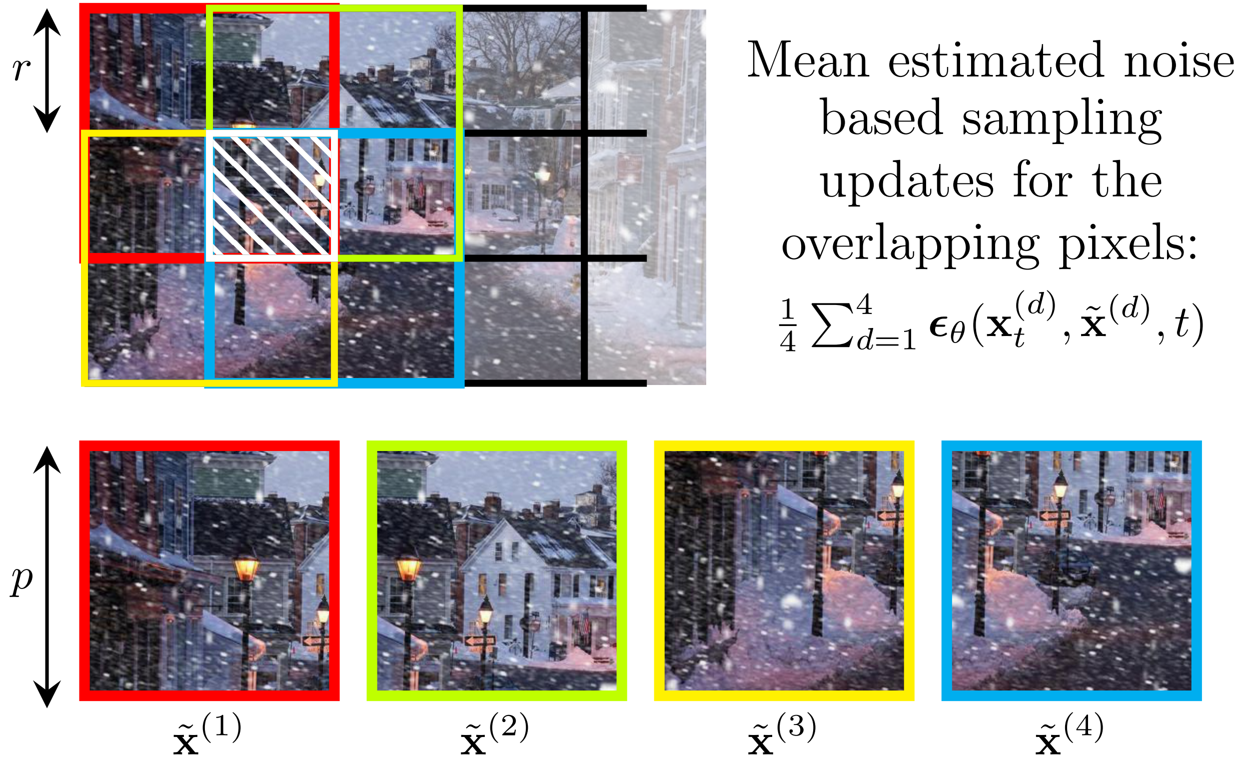

How to merge restored patches into a whole image?

At each sampling time ***t*** = ***T***, … ,***1***

1. use the "noise estimator" network for all overlapping patches to estimate the added noise at time ***t***

2. compute "*mean estimated noise*" based sampling updates for these overlapping regions, and form the restored whole-image at time ***t***

Patch-based Diffusive Image Restoration: Inference

Restoration Process:

Weather-degraded observation: $\tilde{\mathbf{X}}$

D patch pairs

Patch-based Diffusive Image Restoration: Inference

Restoration Process:

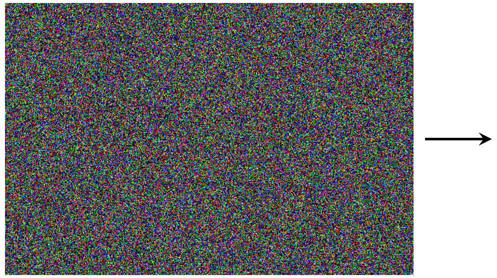

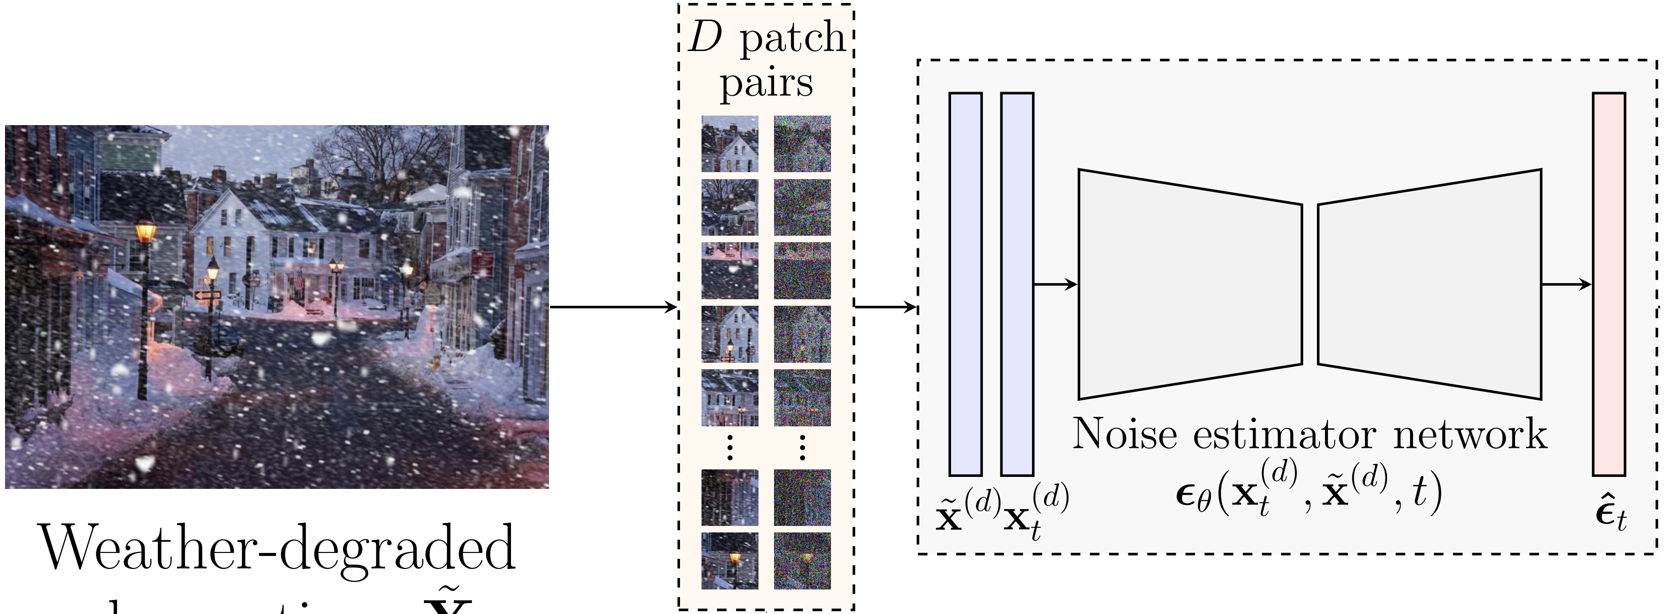

Weather-degraded observation: $\tilde{\mathbf{X}}$

Patch-based Diffusive Image Restoration: Inference

Restoration Process:

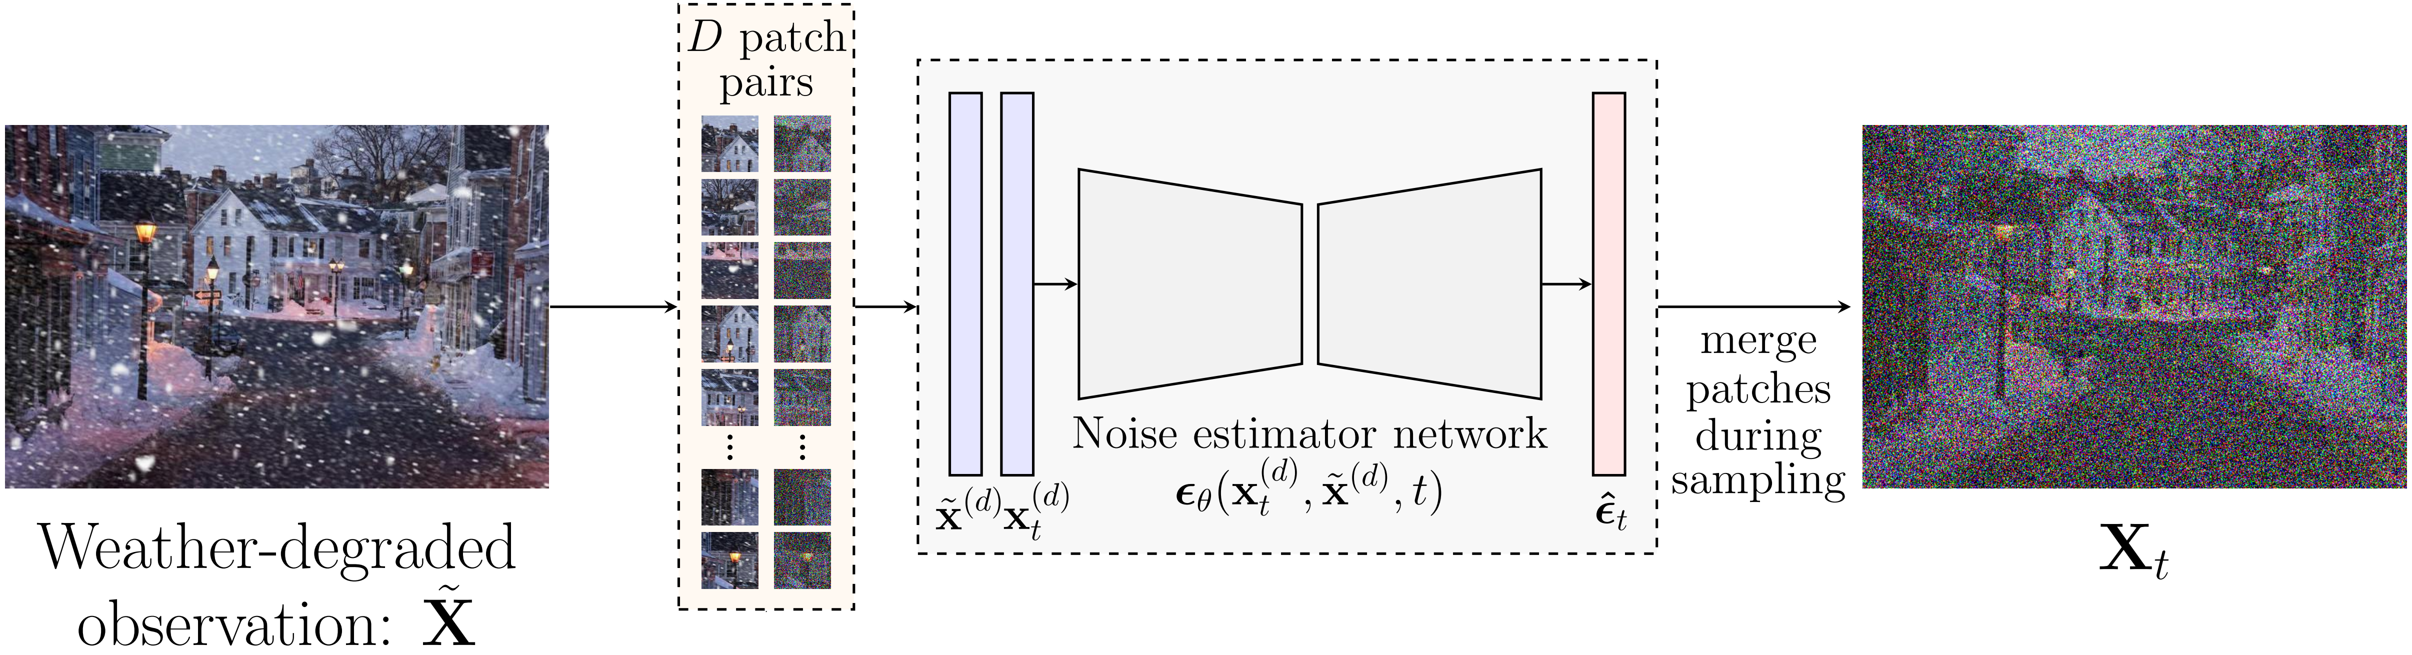

Patch-based Diffusive Image Restoration: Inference

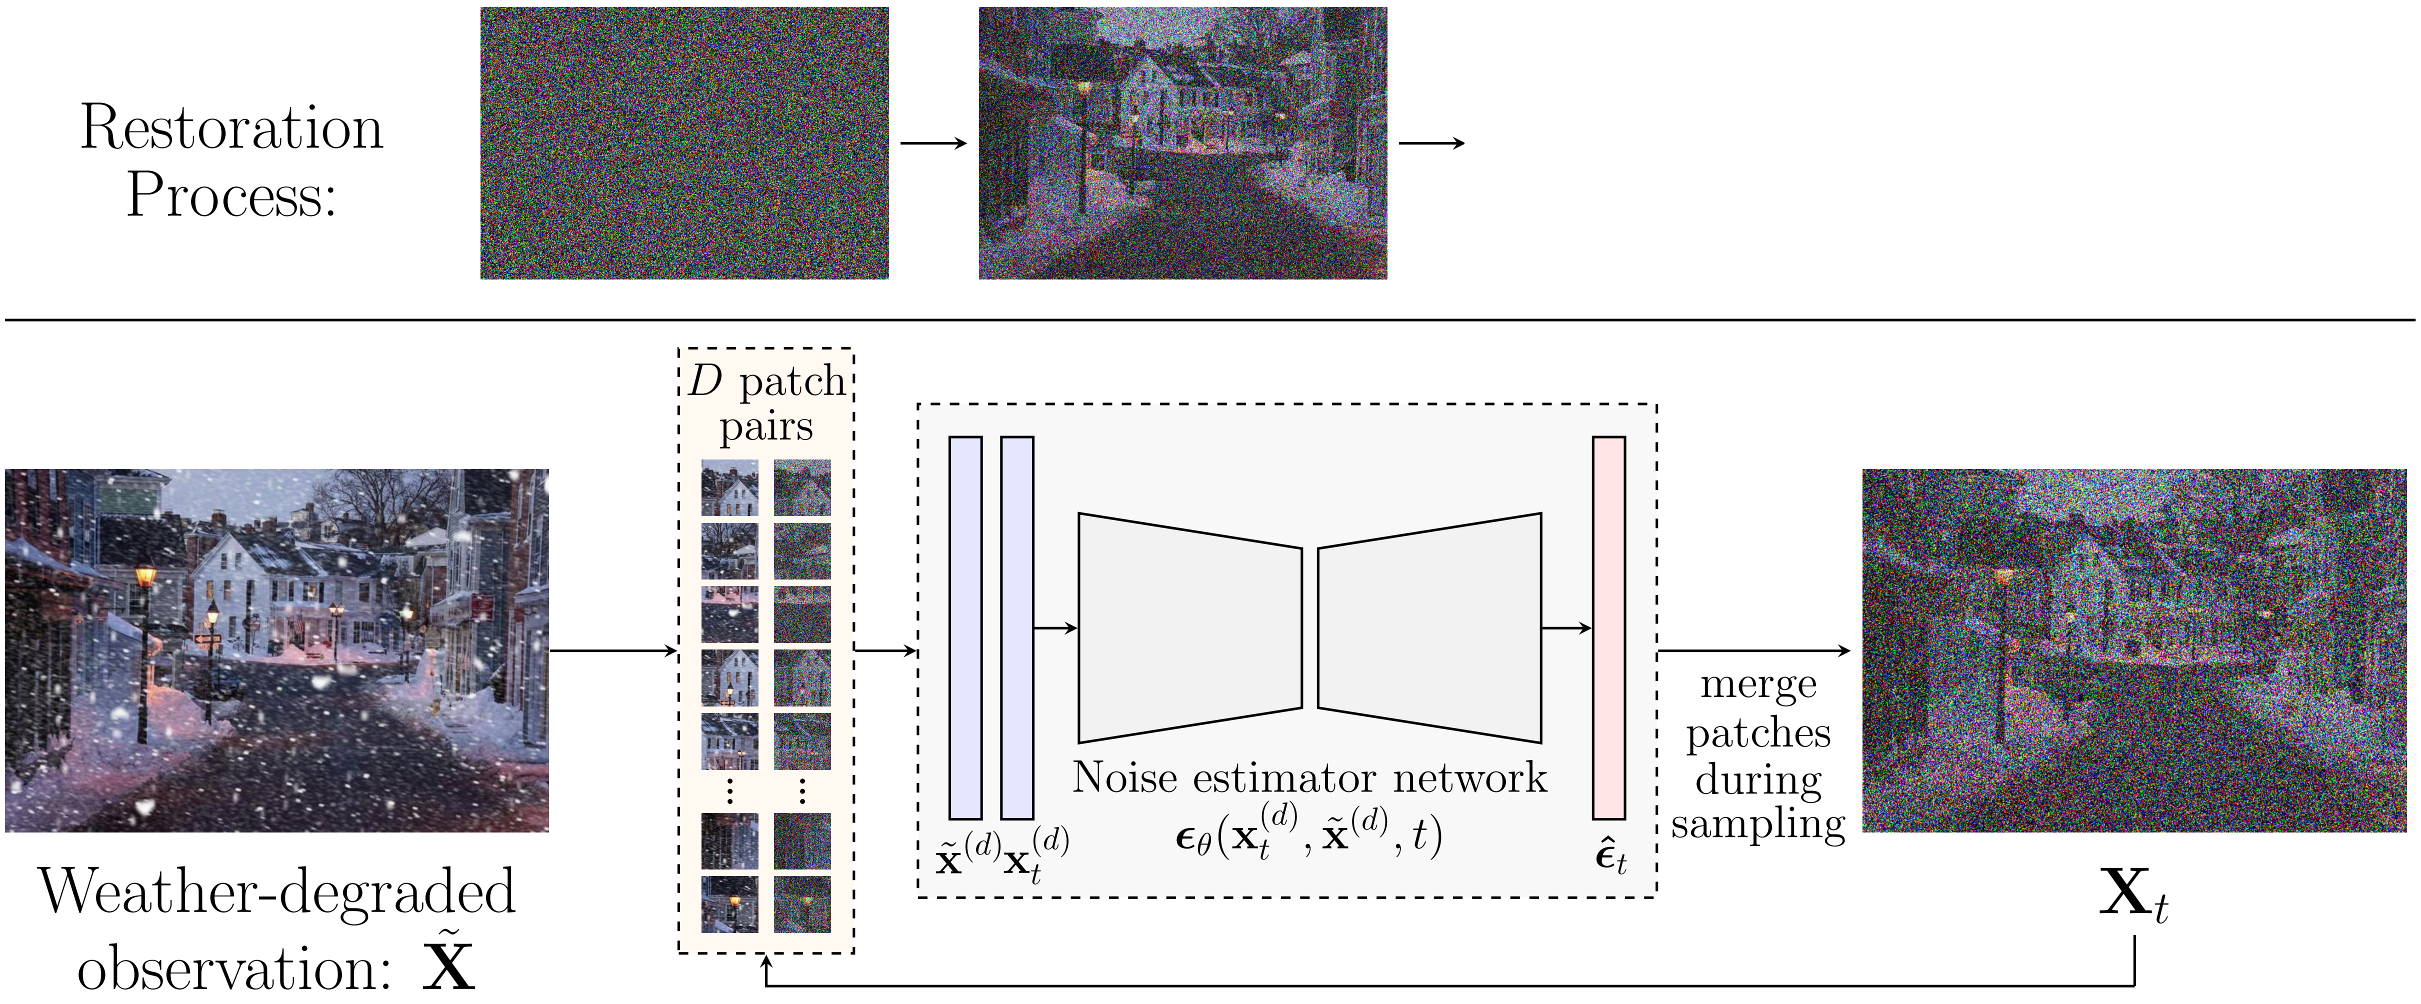

Patch-based Diffusive Image Restoration: Inference

Restoration Process:

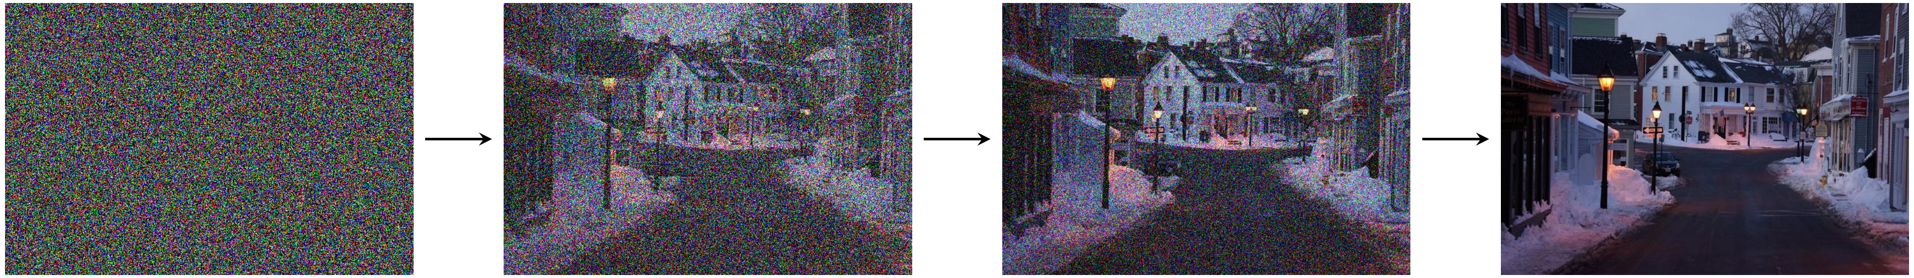

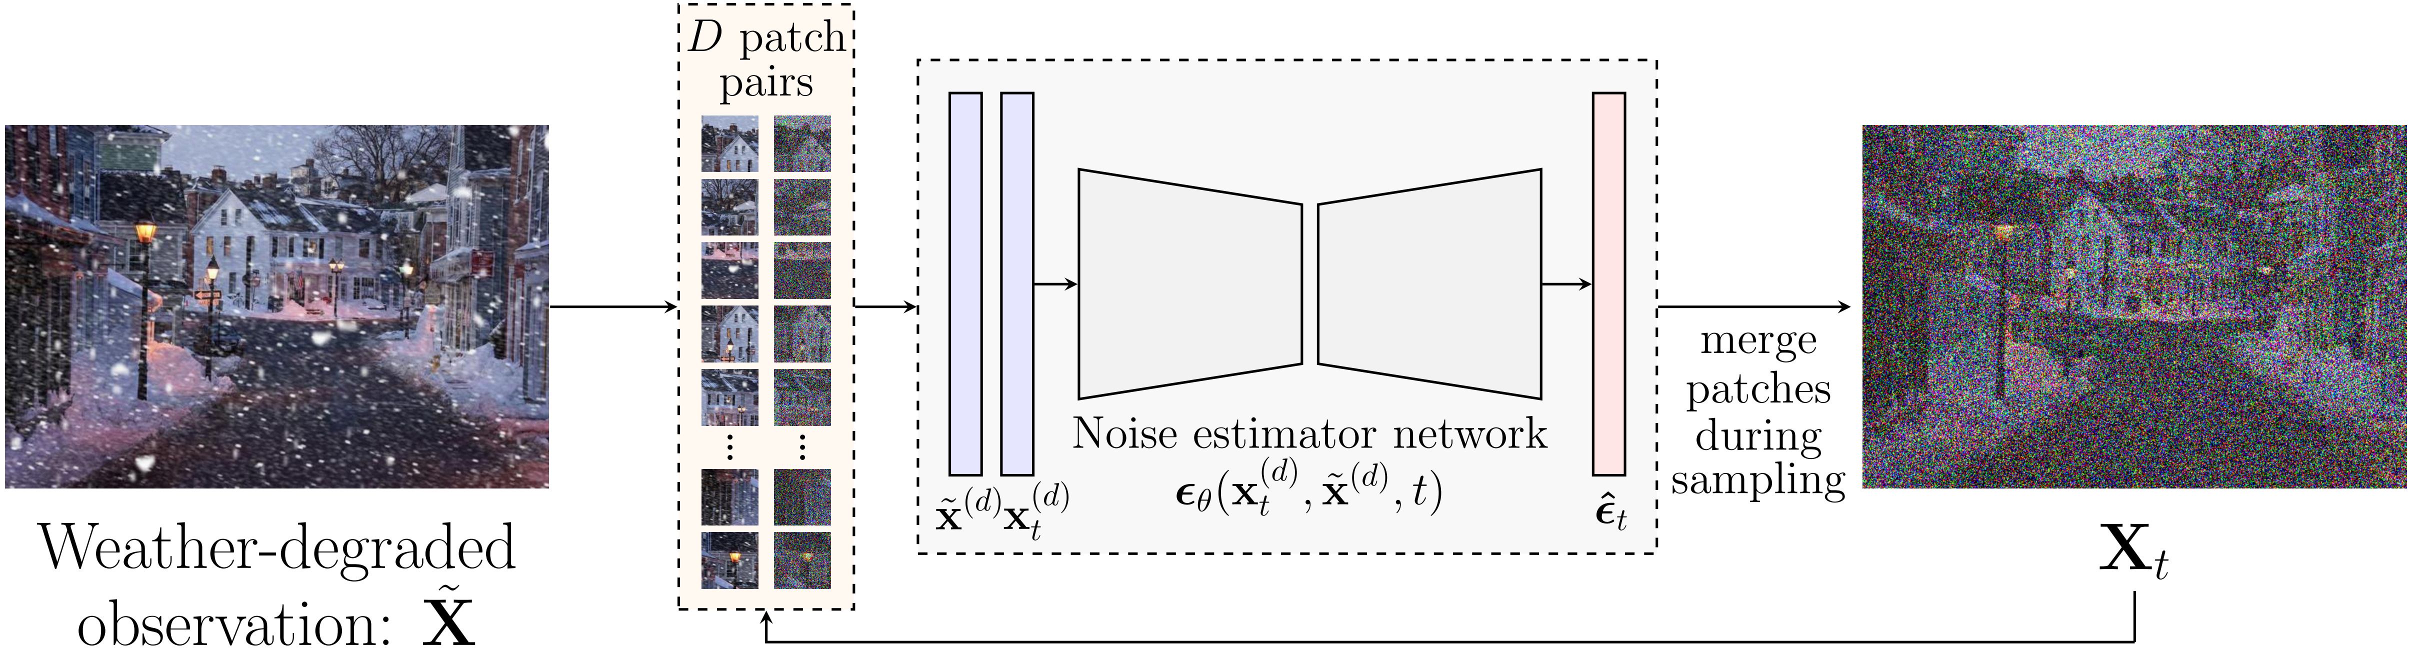

Patch-based Diffusive Image Restoration: Examples

Patch-based Diffusive Image Restoration: Examples

Patch-based Diffusive Image Restoration: Examples

Patch-based Diffusive Image Restoration: Examples

Results: Image Desnowing

	Snow100K-S [3]		Snow100K-L [3]	
	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑
SPANet [44]	29.92	0.8260	23.70	0.7930
JSTASR [58]	31.40	0.9012	25.32	0.8076
RESCAN [43]	31.51	0.9032	26.08	0.8108
DesnowNet [3]	32.33	0.9500	27.17	0.8983
DDMSNet [59]	34.34	0.9445	28.85	0.8772
SnowDiff$_{64}$	**36.59**	**0.9626**	**30.43**	**0.9145**
SnowDiff$_{128}$	<u>36.09</u>	<u>0.9545</u>	<u>30.28</u>	<u>0.9000</u>

(a) Image Desnowing

	Outdoor-Rain [14]	
	PSNR ↑	SSIM ↑

(b) Image Deraining & Dehazing

	RainDrop [12]	
	PSNR ↑	SSIM ↑

(c) Removing Raindrops

Fig. 3. Quantitative comparisons in terms of PSNR and SSIM (higher is better) with state-of-the-art image desnowing and deraining methods. Above half of the tables show comparisons of our weather-specific SnowDiff$_p$, RainHazeDiff$_p$ and RainDropDiff$_p$ models individually evaluated for each task. Bottom half of the tables show evaluations of our unified multi-weather model WeatherDiff$_p$ on all three test sets with respect to All-in-One [9] and TransWeather [7] multi-weather restoration methods. Best and second best values are indicated with bold text and underlined text respectively.

Results: Image Desnowing

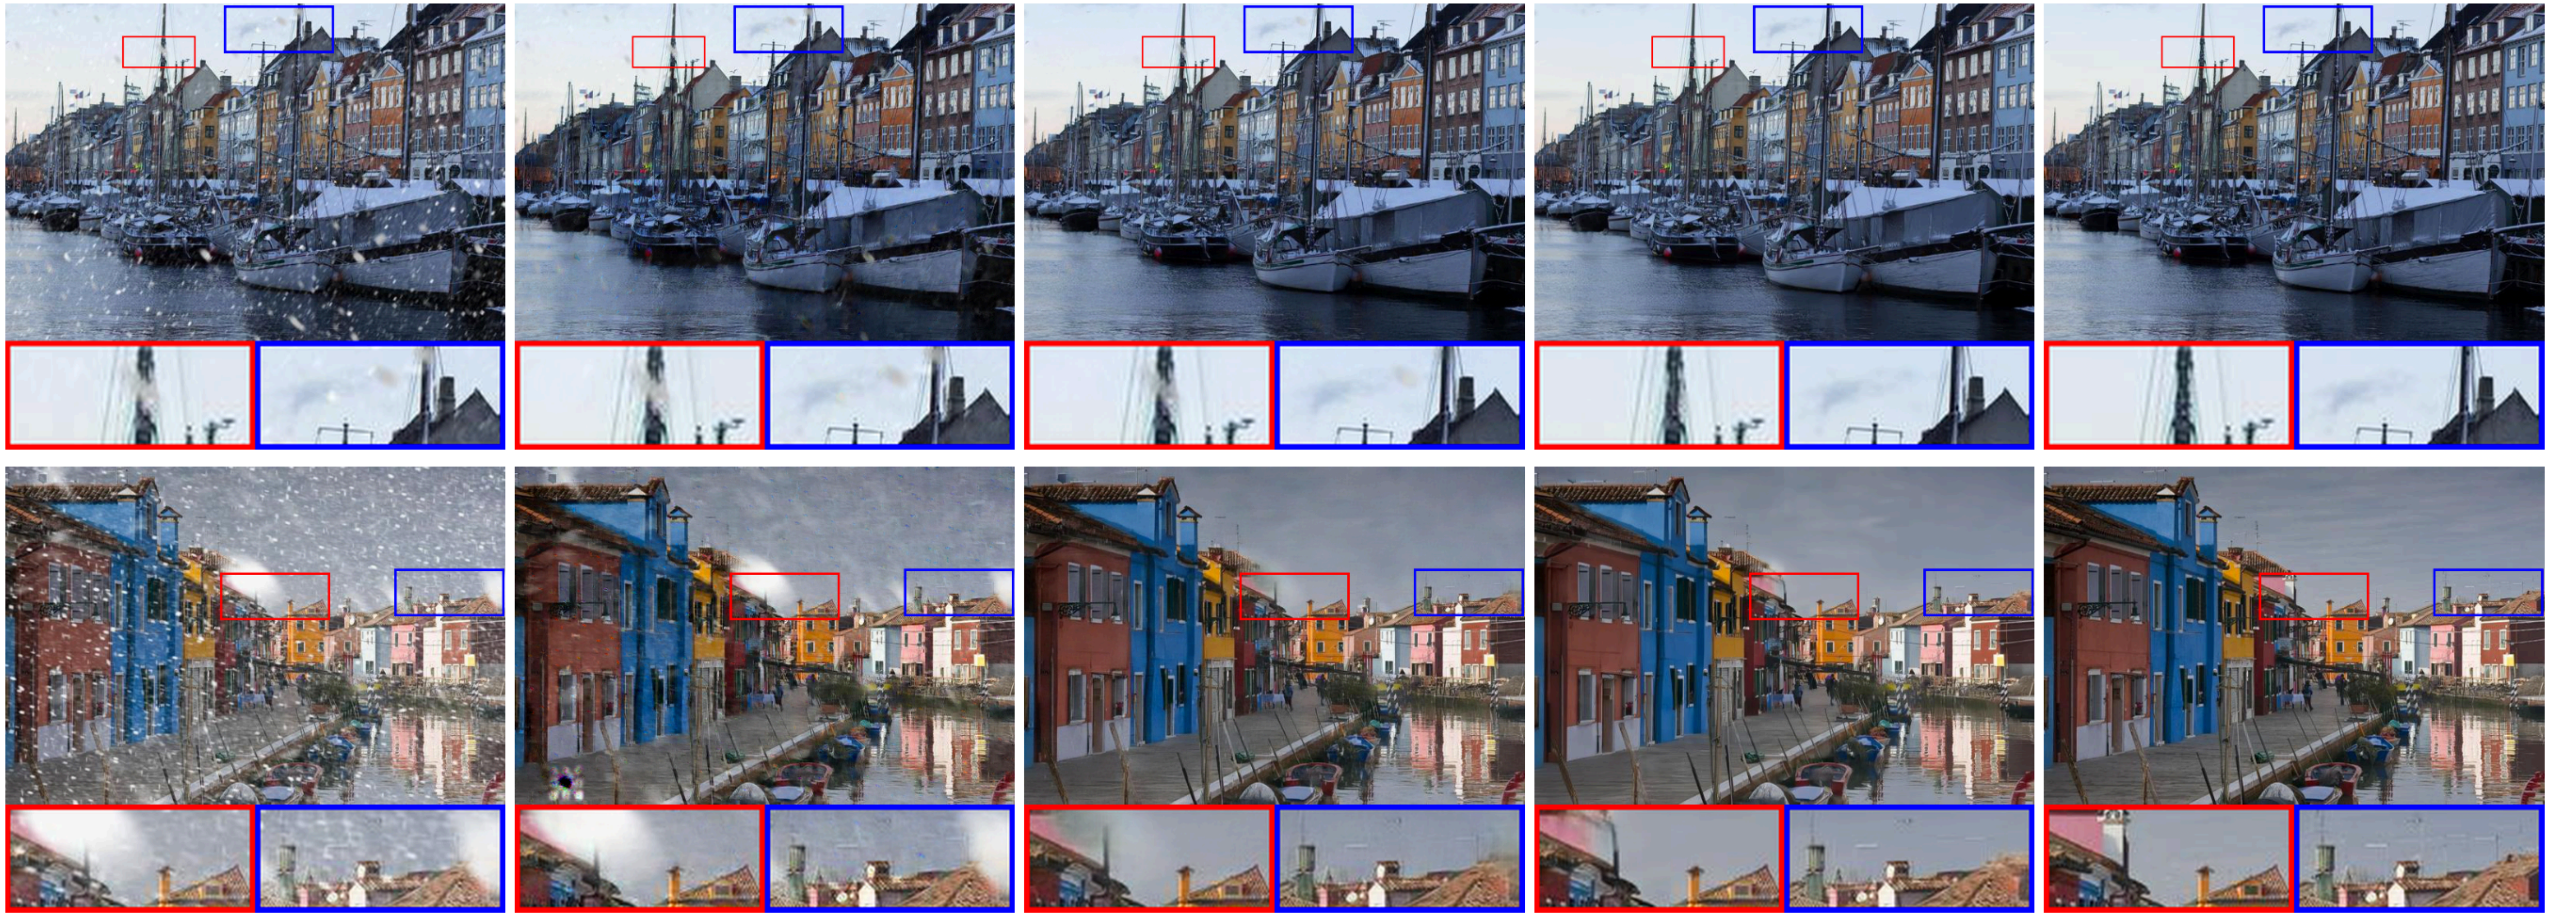

(a) Input (b) DesnowNet [3] (c) DDMSNet [59] (d) **Ours (SnowDiff$_{64}$)** (e) Ground truth

Results: Image Deraining & Dehazing

	Snow100K-S [3]		Snow100K-L [3]	
	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑
SPANet [44]	29.92	0.8260	23.70	0.7930
JSTASR [58]	31.40	0.9012	25.32	0.8076
RESCAN [43]	31.51	0.9032	26.08	0.8108
DesnowNet [3]	32.33	0.9500	27.17	0.8983
DDMSNet [59]	34.34	0.9445	28.85	0.8772
SnowDiff$_{64}$	**36.59**	**0.9626**	**30.43**	**0.9145**
SnowDiff$_{128}$	<u>36.09</u>	<u>0.9545</u>	<u>30.28</u>	<u>0.9000</u>

(a) Image Desnowing

	Outdoor-Rain [14]	
	PSNR ↑	SSIM ↑
CycleGAN [46]	17.62	0.6560
pix2pix [45]	19.09	0.7100
HRGAN [14]	21.56	0.8550
PCNet [53]	26.19	0.9015
MPRNet [54]	<u>28.03</u>	<u>0.9192</u>
RainHazeDiff$_{64}$	**28.38**	**0.9320**
RainHazeDiff$_{128}$	26.84	0.9152

(b) Image Deraining & Dehazing

	RainDrop [12]	
	PSNR ↑	SSIM ↑

(c) Removing Raindrops

Fig. 3. Quantitative comparisons in terms of PSNR and SSIM (higher is better) with state-of-the-art image desnowing and deraining methods. Above half of the tables show comparisons of our weather-specific SnowDiff$_p$, RainHazeDiff$_p$ and RainDropDiff$_p$ models individually evaluated for each task. Bottom half of the tables show evaluations of our unified multi-weather model WeatherDiff$_p$ on all three test sets with respect to All-in-One [9] and TransWeather [7] multi-weather restoration methods. Best and second best values are indicated with bold text and underlined text respectively.

Results: Image Deraining & Dehazing

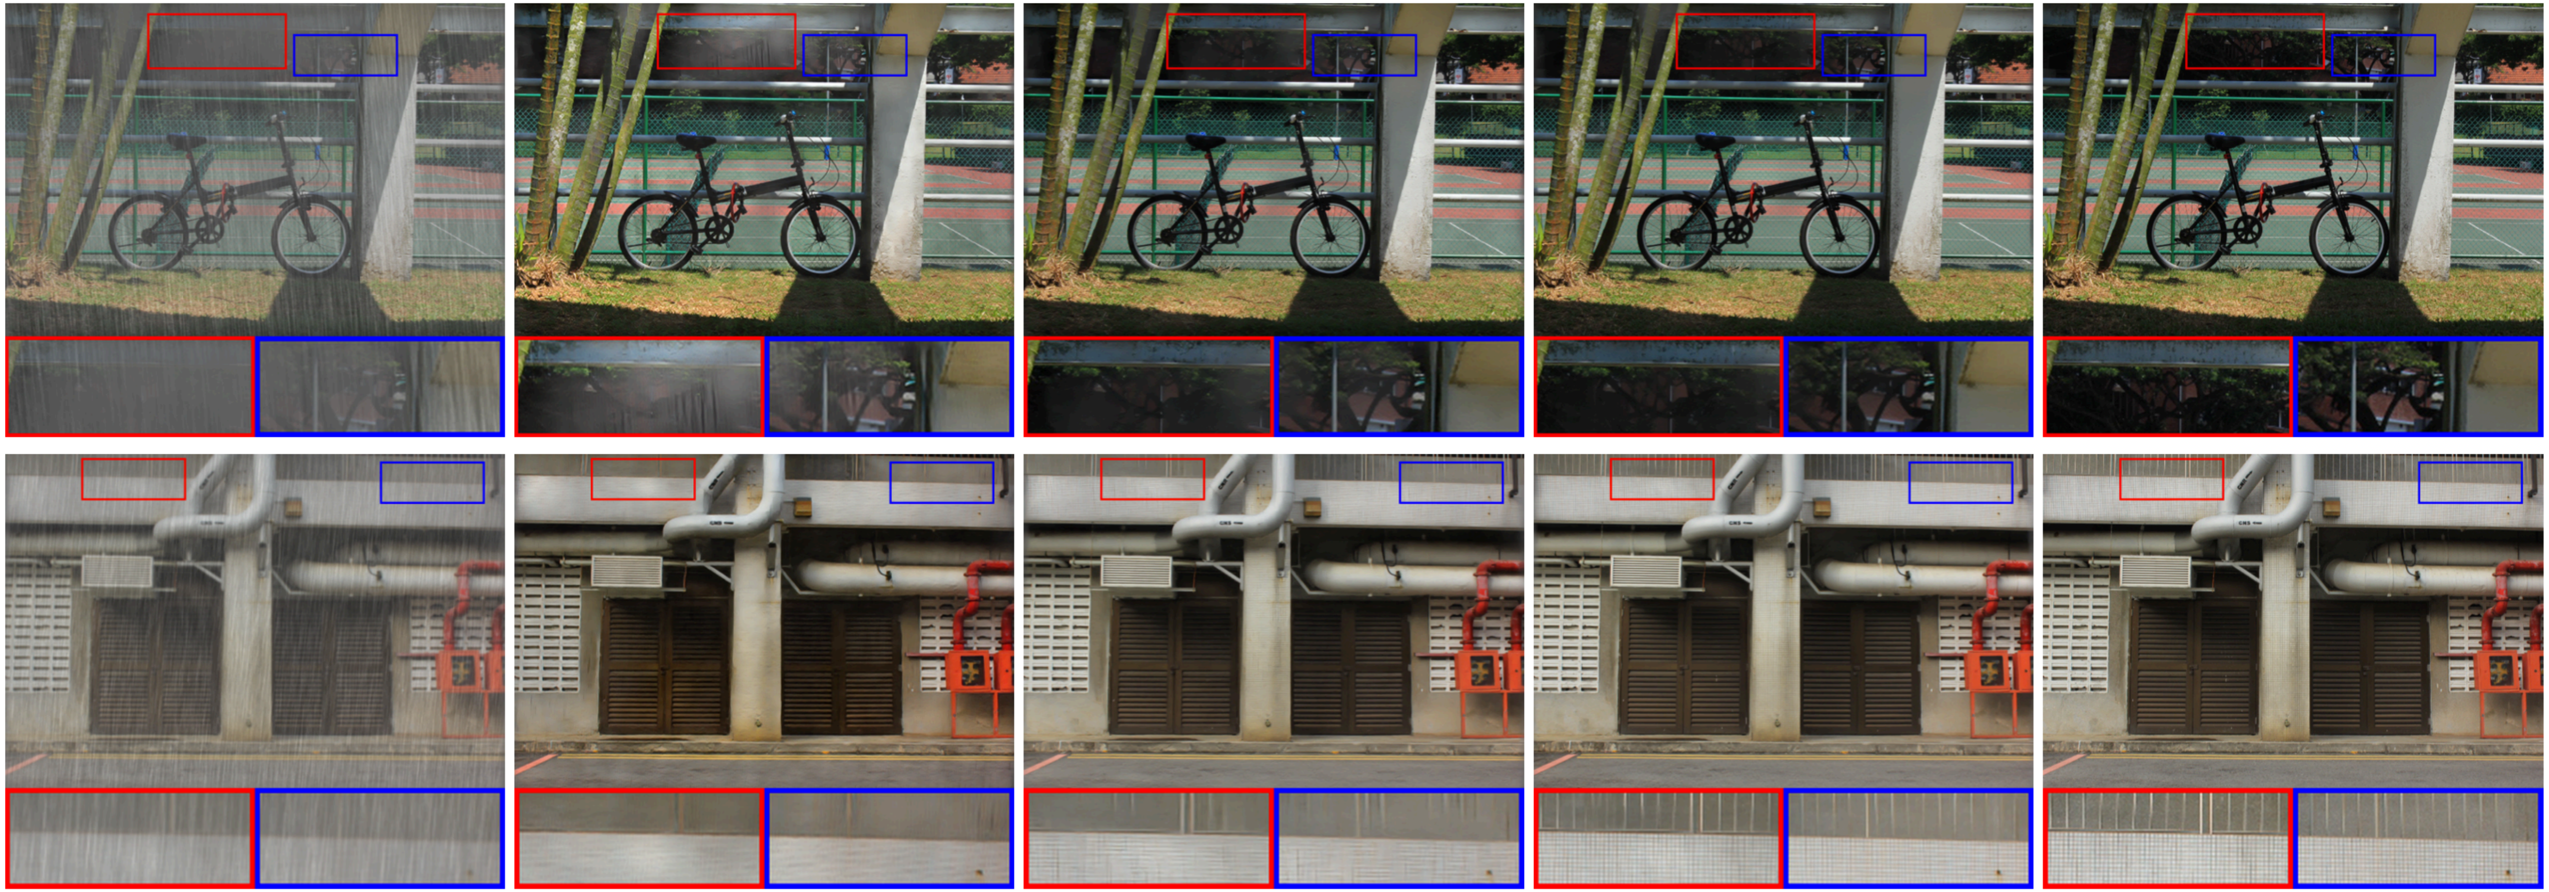

(a) Input (b) HRGAN [14] (c) MPRNet [54] (d) **Ours (RainHazeDiff$_{64}$)** (e) Ground truth

Results: Removing Raindrops

	Snow100K-S [3]		Snow100K-L [3]	
	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑
SPANet [44]	29.92	0.8260	23.70	0.7930
JSTASR [58]	31.40	0.9012	25.32	0.8076
RESCAN [43]	31.51	0.9032	26.08	0.8108
DesnowNet [3]	32.33	0.9500	27.17	0.8983
DDMSNet [59]	34.34	0.9445	28.85	0.8772
SnowDiff$_{64}$	**36.59**	**0.9626**	**30.43**	**0.9145**
SnowDiff$_{128}$	36.09	0.9545	30.28	0.9000

(a) Image Desnowing

	Outdoor-Rain [14]	
	PSNR ↑	SSIM ↑
CycleGAN [46]	17.62	0.6560
pix2pix [45]	19.09	0.7100
HRGAN [14]	21.56	0.8550
PCNet [53]	26.19	0.9015
MPRNet [54]	28.03	0.9192
RainHazeDiff$_{64}$	**28.38**	**0.9320**
RainHazeDiff$_{128}$	26.84	0.9152

(b) Image Deraining & Dehazing

	RainDrop [12]	
	PSNR ↑	SSIM ↑
pix2pix [45]	28.02	0.8547
DuRN [56]	31.24	0.9259
RaindropAttn [55]	31.44	0.9263
AttentiveGAN [12]	31.59	0.9170
IDT [6]	31.87	0.9313
RainDropDiff$_{64}$	32.29	**0.9422**
RainDropDiff$_{128}$	**32.43**	0.9334

(c) Removing Raindrops

Fig. 3. Quantitative comparisons in terms of PSNR and SSIM (higher is better) with state-of-the-art image desnowing and deraining methods. Above half of the tables show comparisons of our weather-specific SnowDiff$_p$, RainHazeDiff$_p$ and RainDropDiff$_p$ models individually evaluated for each task. Bottom half of the tables show evaluations of our unified multi-weather model WeatherDiff$_p$ on all three test sets with respect to All-in-One [9] and TransWeather [7] multi-weather restoration methods. Best and second best values are indicated with bold text and underlined text respectively.

Results: Removing Raindrops

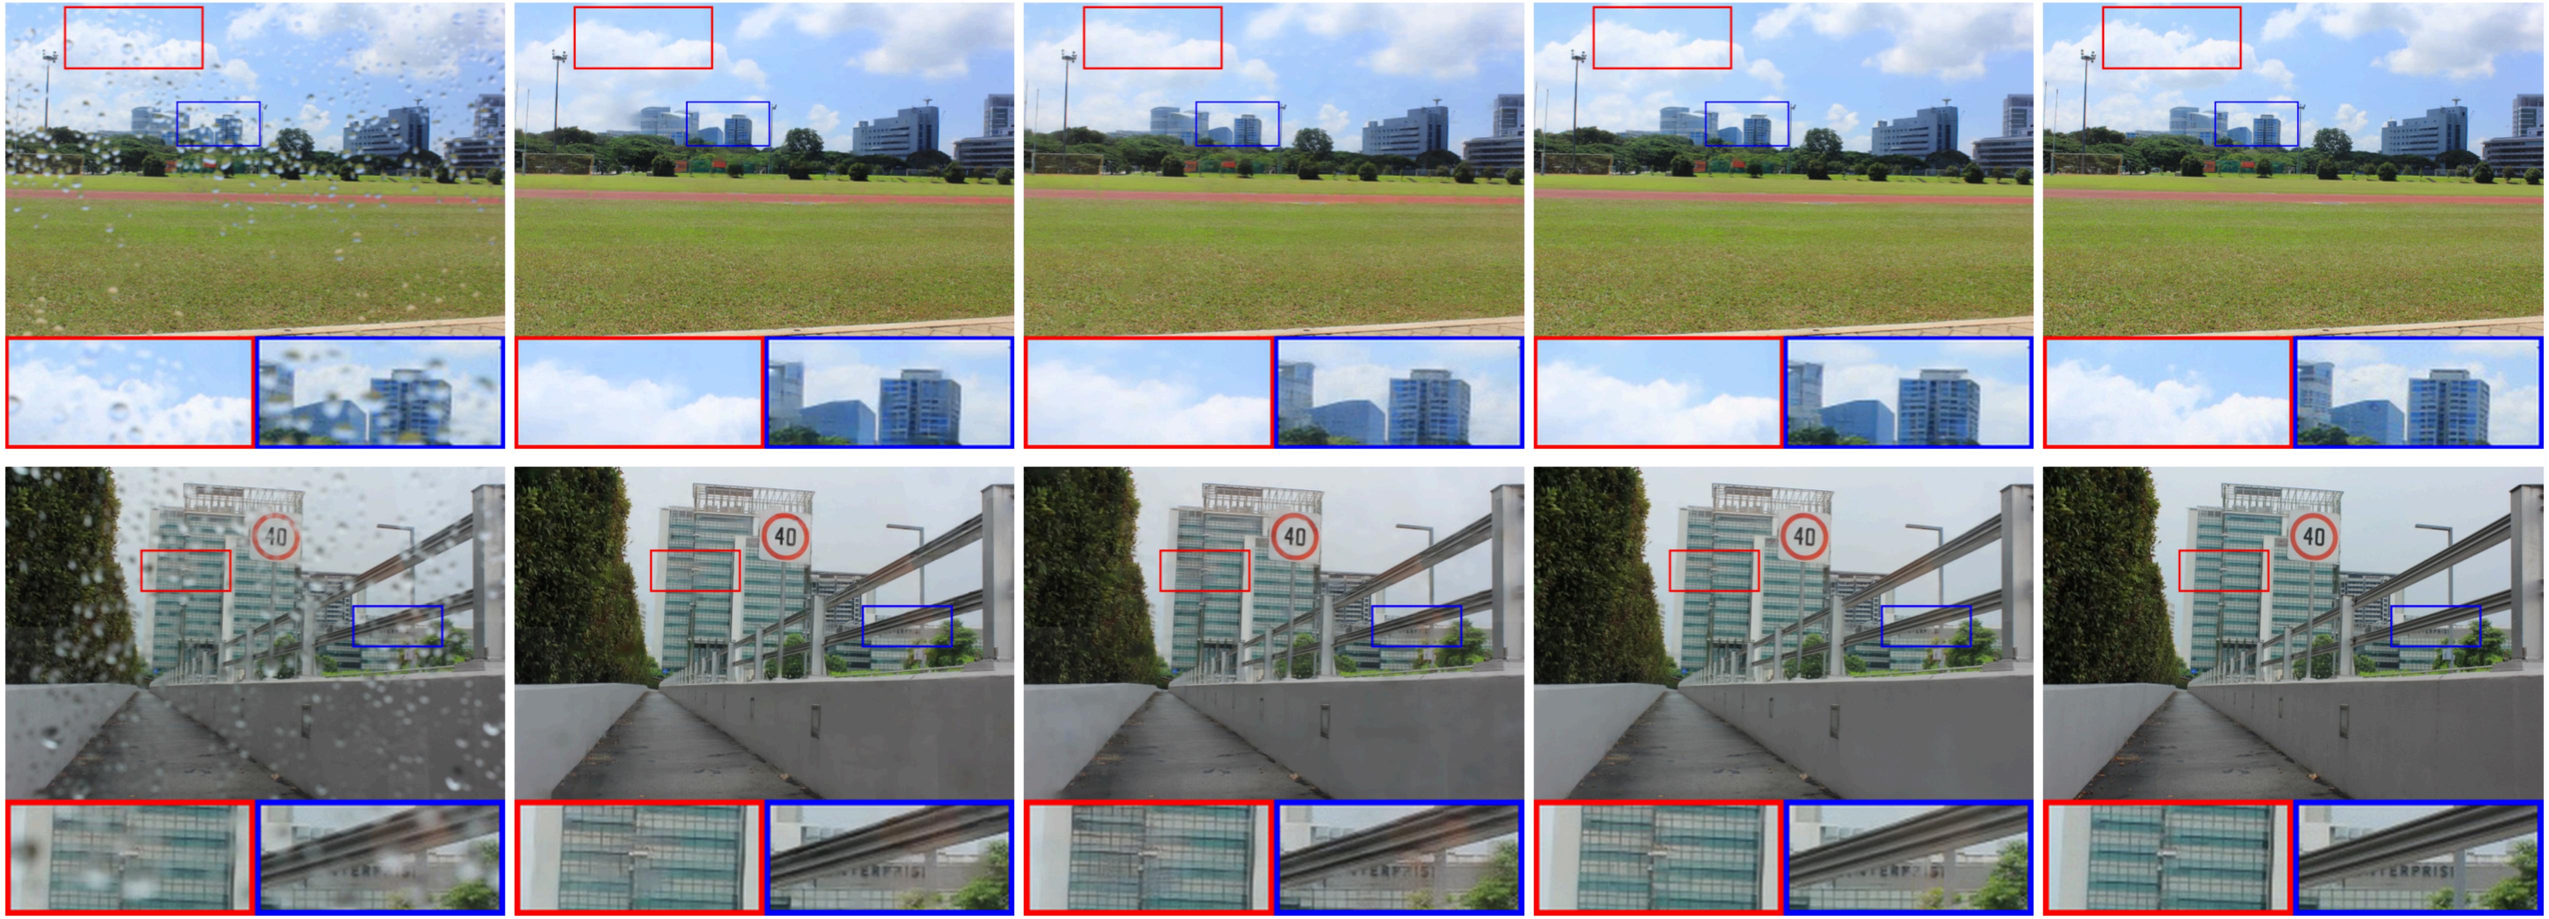

(a) Input (b) RaindropAttn [55] (c) AttentiveGAN [12] (d) **Ours (RainDropDiff$_{128}$)** (e) Ground truth

Results: Multi-Weather Restoration

	Snow100K-S [3]		Snow100K-L [3]	
	PSNR ↑	SSIM ↑	PSNR ↑	SSIM ↑
SPANet [44]	29.92	0.8260	23.70	0.7930
JSTASR [58]	31.40	0.9012	25.32	0.8076
RESCAN [43]	31.51	0.9032	26.08	0.8108
DesnowNet [3]	32.33	0.9500	27.17	0.8983
DDMSNet [59]	34.34	0.9445	28.85	0.8772
SnowDiff$_{64}$	**36.59**	**0.9626**	**30.43**	**0.9145**
SnowDiff$_{128}$	36.09	0.9545	30.28	0.9000
All-in-One [9]	-	-	28.33	0.8820
TransWeather [7]	32.51	0.9341	29.31	0.8879
WeatherDiff$_{64}$	**35.12**	**0.9539**	**29.55**	**0.8988**
WeatherDiff$_{128}$	34.72	0.9509	29.21	0.8911

(a) Image Desnowing

	Outdoor-Rain [14]	
	PSNR ↑	SSIM ↑
CycleGAN [46]	17.62	0.6560
pix2pix [45]	19.09	0.7100
HRGAN [14]	21.56	0.8550
PCNet [53]	26.19	0.9015
MPRNet [54]	28.03	0.9192
RainHazeDiff$_{64}$	**28.38**	**0.9320**
RainHazeDiff$_{128}$	26.84	0.9152
All-in-One [9]	24.71	0.8980
TransWeather [7]	28.83	0.9000
WeatherDiff$_{64}$	28.86	**0.9257**
WeatherDiff$_{128}$	**29.53**	0.9208

(b) Image Deraining & Dehazing

	RainDrop [12]	
	PSNR ↑	SSIM ↑
pix2pix [45]	28.02	0.8547
DuRN [56]	31.24	0.9259
RaindropAttn [55]	31.44	0.9263
AttentiveGAN [12]	31.59	0.9170
IDT [6]	31.87	0.9313
RainDropDiff$_{64}$	32.29	**0.9422**
RainDropDiff$_{128}$	**32.43**	0.9334
All-in-One [9]	**31.12**	0.9268
TransWeather [7]	30.17	0.9157
WeatherDiff$_{64}$	30.26	**0.9277**
WeatherDiff$_{128}$	29.37	0.9213

(c) Removing Raindrops

Fig. 3. Quantitative comparisons in terms of PSNR and SSIM (higher is better) with state-of-the-art image desnowing and deraining methods. Above half of the tables show comparisons of our weather-specific SnowDiff$_p$, RainHazeDiff$_p$ and RainDropDiff$_p$ models individually evaluated for each task. Bottom half of the tables show evaluations of our unified multi-weather model WeatherDiff$_p$ on all three test sets with respect to All-in-One [9] and TransWeather [7] multi-weather restoration methods. Best and second best values are indicated with bold text and underlined text respectively.

Results: Real-World Image Restoration

Input Image | TransWeather | **<u>Ours (WeatherDiff)</u>**

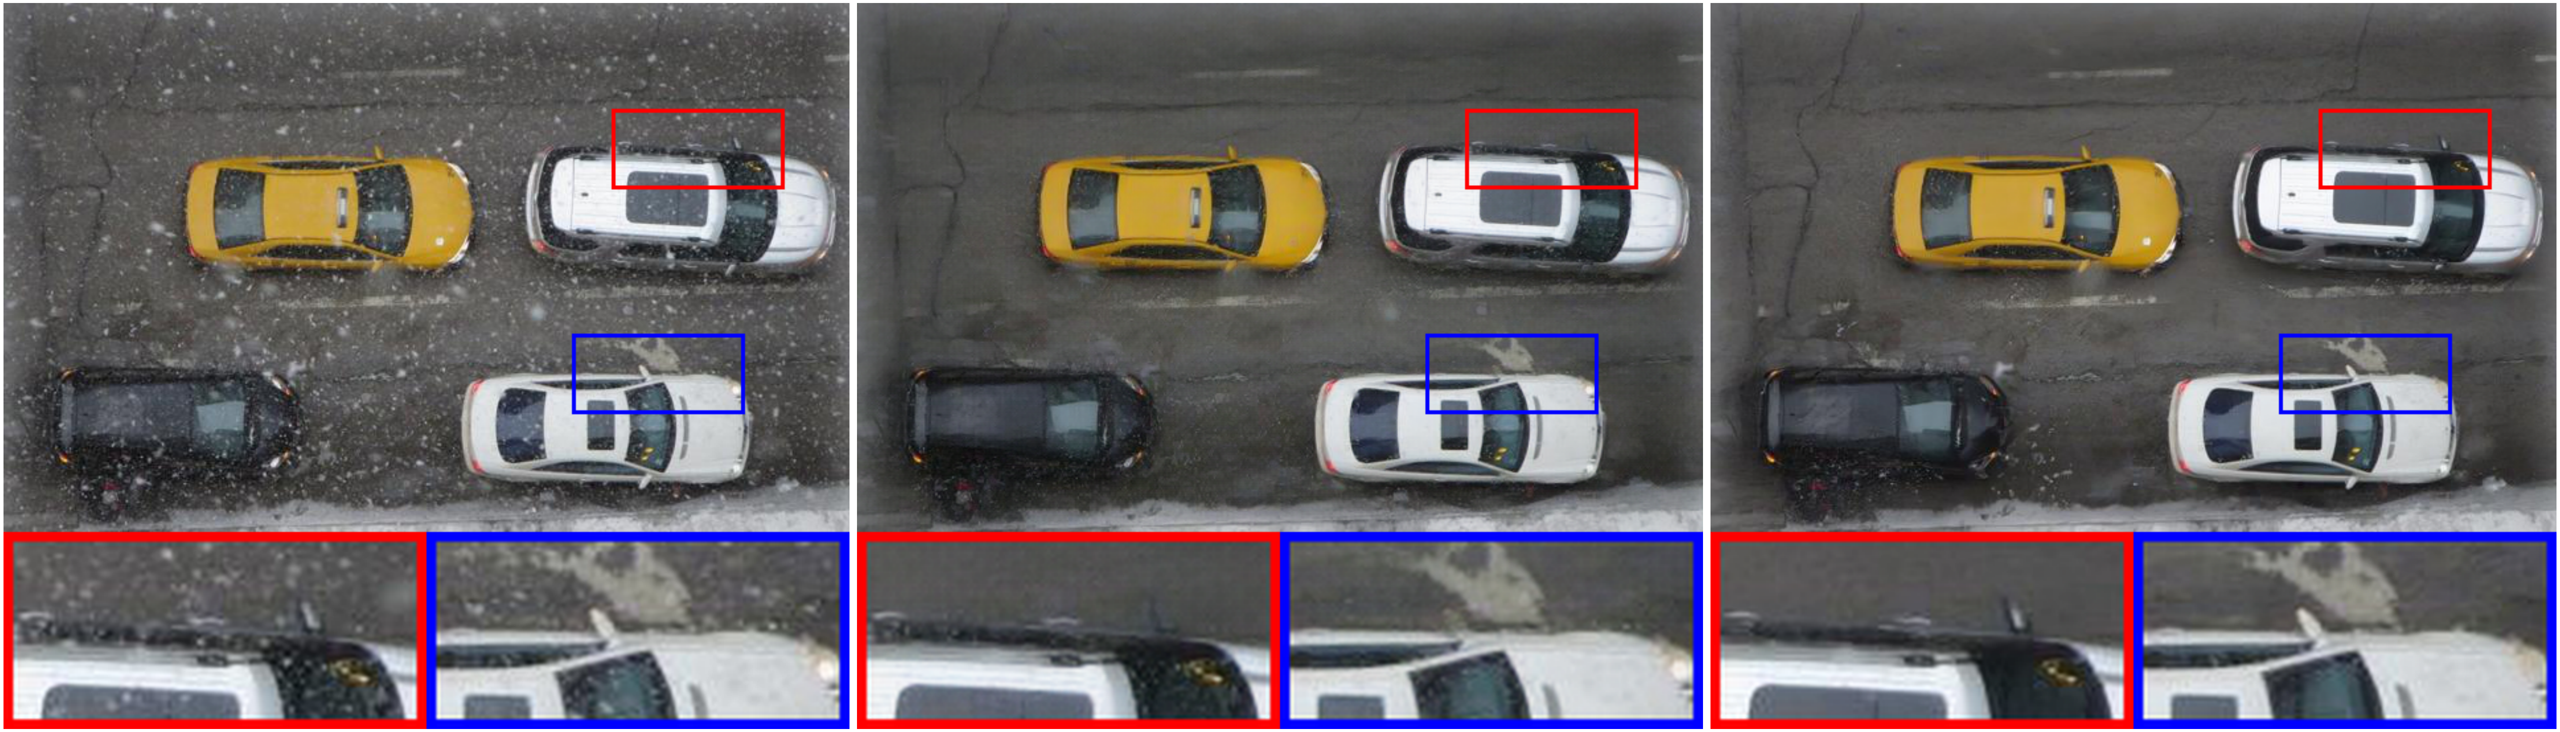

Results: Real-World Image Restoration

Input Image | TransWeather | **Ours (WeatherDiff)**

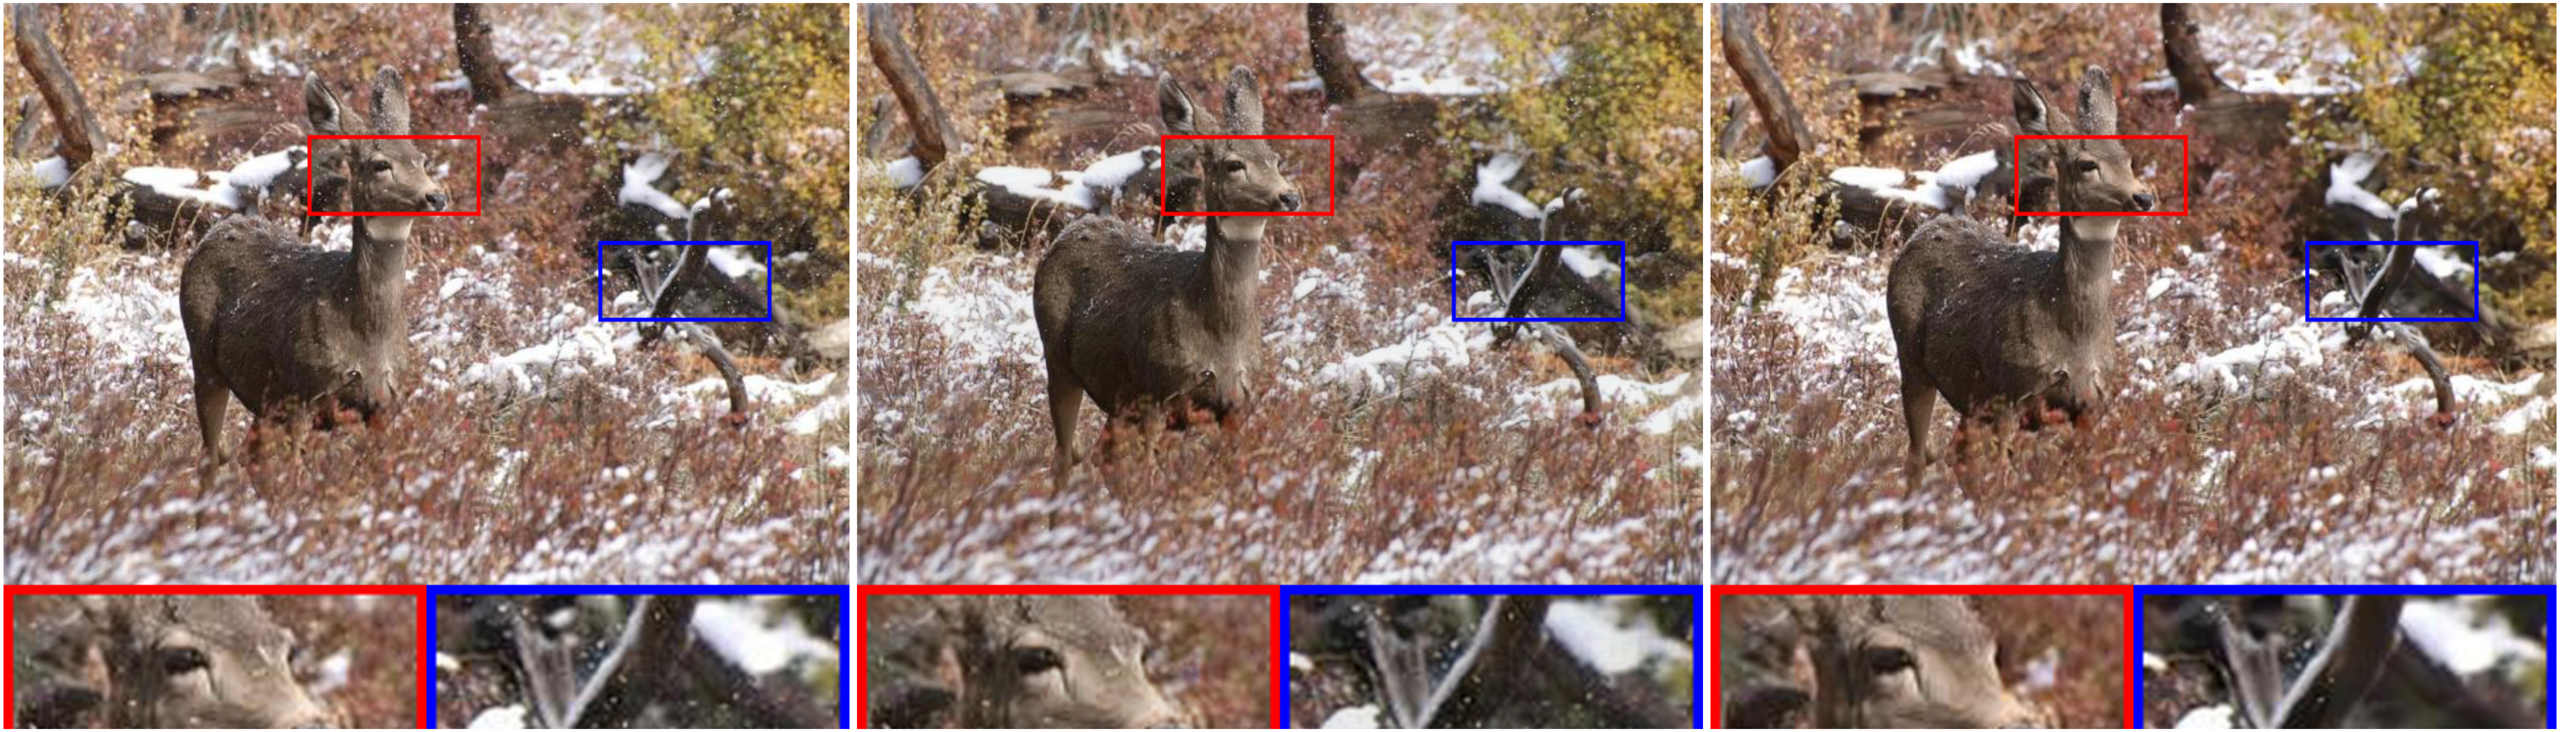

Results: Real-World Image Restoration

Input Image

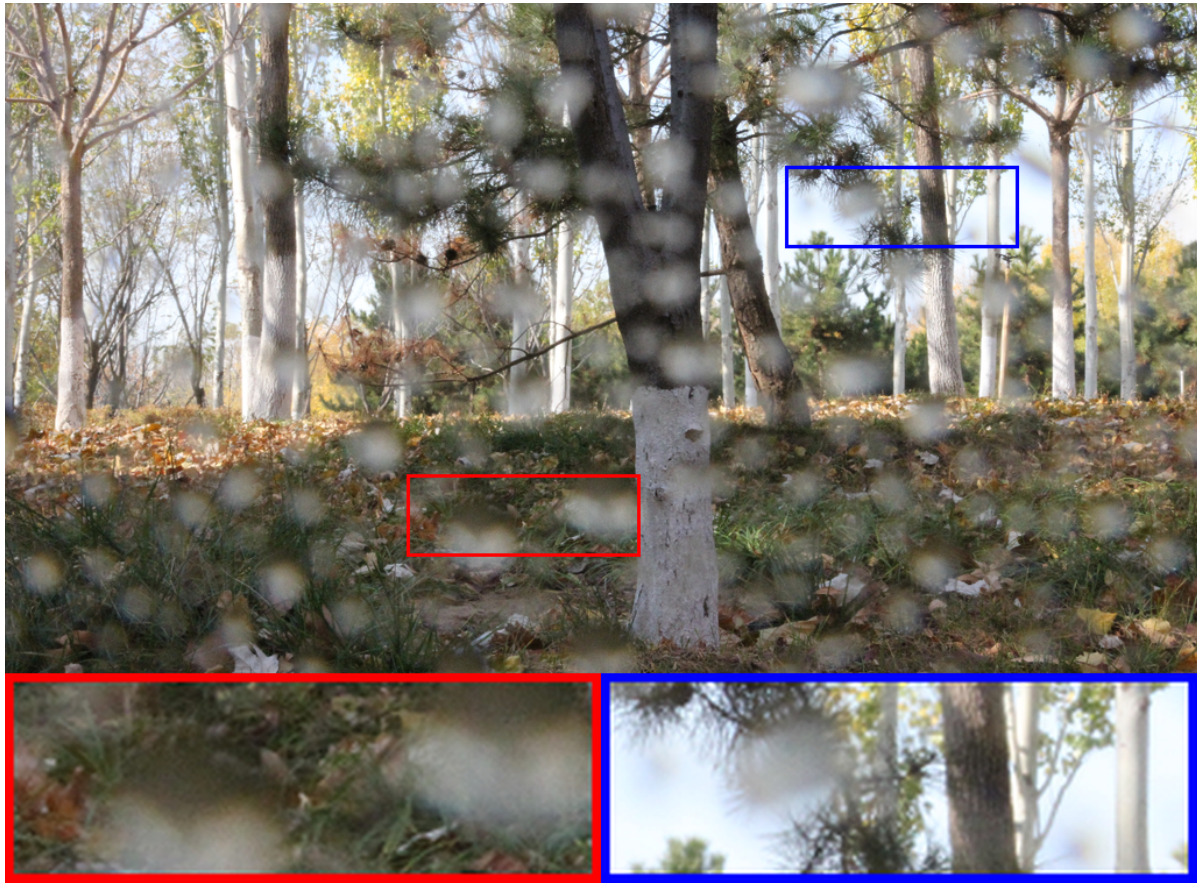

Ours (WeatherDiff)

Results: Real-World Image Restoration

Thank you for your attention!

Paper: https://arxiv.org/pdf/2207.14626.pdf

Code: https://github.com/IGITUGraz/WeatherDiffusion

TU Graz

SAL

SILICON AUSTRIA LABS

Restoring Vision in Adverse Weather Conditions with Patch-Based Denoising Diffusion Models

Ozan Özdenizci and Robert Legenstein

Abstract—Image restoration under adverse weather conditions has been of significant interest for various computer vision applications. Recent successful methods rely on the current progress in deep neural network architectural designs (e.g., with vision transformers). Motivated by the recent progress achieved with state-of-the-art conditional generative models, we present a novel patch-based image restoration algorithm based on denoising diffusion probabilistic models. Our patch-based diffusion modeling approach enables size-agnostic image restoration by using a guided denoising process with smoothed noise estimates across overlapping patches during inference. We empirically evaluate our model on benchmark datasets for image desnowing, combined deraining and dehazing, and raindrop removal. We demonstrate our approach to achieve state-of-the-art performances on both weather-specific and multi-weather image restoration, and qualitatively show strong generalization to real-world test images.

Index Terms—denoising diffusion models, patch-based image restoration, deraining, desnowing, dehazing, raindrop removal.

1 Introduction

The restoration of images under adverse impacts of weather conditions such as heavy rain or snow is of wide interest to computer vision research. At the extreme, observed images to be restored may contain severe weather related obstructions of the true background (e.g., snow flakes, dense hazing effects), causing a well known ill-posed inverse problem where various solutions can be obtained for the unknown ground truth background. Deep neural networks (DNNs) are shown to excel at such image restoration tasks compared to traditional approaches [1], [2], [3], and this success extends with the current progress in DNN architectural designs, e.g., with vision transformers [4], [5]. State-of-the-art designs have recently shown its effectiveness in low-level weather restoration problems with transformers [6], [7] and multi-layer perceptron based models [8]. Beyond task-specialized solutions, recent work also proposed to tackle this problem for multiple weather corruptions in unified architectures [7], [9], [10], [11].

Earlier deep learning based solutions to adverse weather restoration have extensively explored task-specific generative modeling methods, mainly with generative adversarial networks (GANs) [12], [13], [14]. In this setting generative models aim to learn the underlying data distribution for cleared image backgrounds, given weather-degraded examples from a training set. Due to their stronger expressiveness in that sense, generative approaches further accommodate the potential of better generalization to multi-task vision restoration problems. Along this line, we introduce a novel solution to this problem by using a state-of-the-art conditional generative modeling approach, with denoising diffusion probabilistic models [15], [16].

Denoising diffusion models have recently demonstrated remarkable success in various generative modeling tasks [17], [18], [19], [20]. These architectures were however not yet considered for image restoration under adverse weather conditions, or demonstrated to generalize across multiple image restoration problems. A major obstacle for their usage in image restoration is their architectural constraint that prohibits size-agnostic image restoration, whereas image restoration benchmarks and real-world problems consist of images with various sizes.

We present a novel perspective to the problem of improving vision in adverse weather conditions using denoising diffusion models. Particularly for image restoration, we introduce a novel patch-based diffusive restoration approach to enable size-agnostic processing. Our method uses a guided denoising process for diffusion models by steering the sampling process based on smoothed noise estimates for overlapping patches. Proposed patch-based image processing scheme further introduces a light-weight diffusion modeling approach, and extends practicality of state-of-the-art diffusion models with extensive computational resource demands. We experimentally use extreme weather degradation benchmarks on removing snow, combined rain with haze, and removal of raindrops obstructing the camera sensor. We demonstrate our diffusion modeling perspective to excel at several associated problems.

Our contributions are summarized as follows:

- We present a novel patch-based diffusive image restoration algorithm for arbitrary sized image processing with denoising diffusion models.
- We empirically demonstrate our approach to achieve state-of-the-art performance on both weather-specific and multi-weather restoration tasks.
- We qualitatively present strong generalization from synthetic to real-world multi-weather restoration with our generative modeling perspective.

- *O. Özdenizci and R. Legenstein are with the Institute of Theoretical Computer Science, Graz University of Technology, Graz, Austria. E-mail: {ozan.ozdenizci,robert.legenstein}@igi.tugraz.at*
- *O. Özdenizci is also affiliated with TU Graz - SAL Dependable Embedded Systems Lab, Silicon Austria Labs, Graz, Austria.*