Diffusion Application on low-level vision

What is low-level?

Low-level task aims to convert degraded images to enhanced images, such as Super-Resolution, denoise, deblur, dehaze, low-light enhancement, deartifacts.....

Supervised Method:

Deep neural network Paired data Encoder-Decoder



GAN Method



Palette: Image-to-Image Diffusion Models

Chitwan Saharia, William Chan, Huiwen Chang, Chris A Lee, Jonathan Ho, Tim Salimans David J Fleet, Mohammad Norouzi Google Research, Brain Team Canada {sahariac,williamchan,davidfleet,mnorouzi}@google.com

Abstract

Palette develops a unified framework for image-to-image translation based on conditional diffusion models and evaluates this framework on four challenging image-to-image translation tasks, namely colorization, inpainting, uncropping, and JPEG restoration.

This implementation of image-to-image diffusion models out performs strong GAN and regression baselines on all tasks, without task-specific hyper-parameter tuning, architecture customization, or any auxiliary loss or sophisticated new techniques needed.

Method

x: degraded image

Origin Diffusion

The forward diffusion process is a Markovian process that iteratively adds Gaussia noise over T iterations:

$$q(\boldsymbol{y}_{t+1}|\boldsymbol{y}_t) = \mathcal{N}(\boldsymbol{y}_{t-1}; \sqrt{\alpha_t} \boldsymbol{y}_{t-1}, (1-\alpha_t)I)$$
(2)
$$q(\boldsymbol{y}_{1:T}|\boldsymbol{y}_0) = \prod_{t=1}^T q(\boldsymbol{y}_t|\boldsymbol{y}_{t-1})$$
(3)

Note, we can also marginalize the forward process at each step:

$$q(\boldsymbol{y}_t|\boldsymbol{y}_0) = \mathcal{N}(\boldsymbol{y}_t; \sqrt{\gamma_t} \boldsymbol{y}_0, (1-\gamma_t)I), \qquad (4)$$

where $\gamma_t = \prod_{t'}^t \alpha'_t$.

Use Bayer's rule:

$$q(\boldsymbol{y}_{t-1} \mid \boldsymbol{y}_0, \boldsymbol{y}_t) = \mathcal{N}(\boldsymbol{y}_{t-1} \mid \boldsymbol{\mu}, \sigma^2 \boldsymbol{I})$$
(5)

where
$$\boldsymbol{\mu} = \frac{\sqrt{\gamma_{t-1}} (1-\alpha_t)}{1-\gamma_t} \boldsymbol{y}_0 + \frac{\sqrt{\alpha_t} (1-\gamma_{t-1})}{1-\gamma_t} \boldsymbol{y}_t$$
 and $\sigma^2 = \frac{(1-\gamma_{t-1})(1-\alpha_t)}{1-\gamma_t}$

Palette Training:

Palette learns a reverse process which inverts the forward process. Given a noisy image \tilde{y} ,

$$\widetilde{\boldsymbol{y}} = \sqrt{\gamma} \, \boldsymbol{y}_0 + \sqrt{1 - \gamma} \, \boldsymbol{\epsilon} \,, \, \boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{I}) \,, \tag{6}$$

the goal is to recover the target image y_0

Learning entails prediction of the noise vector ε by optimizing the objective

$$\mathbb{E}_{(\boldsymbol{x},\boldsymbol{y})}\mathbb{E}_{\boldsymbol{\epsilon},\boldsymbol{\gamma}}\left\|f_{\boldsymbol{\theta}}(\boldsymbol{x},\underbrace{\sqrt{\boldsymbol{\gamma}}\,\boldsymbol{y}_{0}+\sqrt{1-\boldsymbol{\gamma}}\,\boldsymbol{\epsilon}}_{\boldsymbol{y}},\boldsymbol{\gamma})-\boldsymbol{\epsilon}\right\|_{p}^{p}.$$

Palette Sample:

given y_t , we approximate y_0 by rearranging terms in equation 6

$$\hat{\boldsymbol{y}}_0 = \frac{1}{\sqrt{\gamma_t}} \left(\boldsymbol{y}_t - \sqrt{1 - \gamma_t} f_{\theta}(\boldsymbol{x}, \boldsymbol{y}_t, \gamma_t) \right) \,.$$

With this parameterization, each iteration of the reverse process can be computed as

$$\boldsymbol{y}_{t-1} \leftarrow \frac{1}{\sqrt{\alpha_t}} \left(\boldsymbol{y}_t - \frac{1 - \alpha_t}{\sqrt{1 - \gamma_t}} f_{\theta}(\boldsymbol{x}, \boldsymbol{y}_t, \gamma_t) \right) + \sqrt{1 - \alpha_t} \boldsymbol{\epsilon}_t ,$$

Palette

Algorithm 1 Training a denoising model f_{θ}

1: repeat

- 2: $(\boldsymbol{x}, \boldsymbol{y}_0) \sim p(\boldsymbol{x}, \boldsymbol{y})$
- 3: $\gamma \sim p(\gamma)$
- 4: $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
- 5: Take a gradient descent step on $\nabla_{\theta} \left\| f_{\theta}(\boldsymbol{x}, \sqrt{\gamma} \boldsymbol{y}_{0} + \sqrt{1 - \gamma} \boldsymbol{\epsilon}, \gamma) - \boldsymbol{\epsilon} \right\|_{p}^{p}$ 6: **until** converged

Algorithm 2 Inference in *T* iterative refinement steps

1:
$$y_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

2: for $t = T, ..., 1$ do
3: $z \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $t > 1$, else $z = \mathbf{0}$
4: $y_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(y_t - \frac{1-\alpha_t}{\sqrt{1-\gamma_t}} f_\theta(\mathbf{x}, \mathbf{y}_t, \gamma_t) \right) + \sqrt{1-\alpha_t} z$
5: end for
6: return y_0

DDPM

Algorithm 1 Training

- 1: repeat
- 2: $\mathbf{x}_0 \sim q(\mathbf{x}_0)$
- 3: $t \sim \text{Uniform}(\{1, \dots, T\})$
- 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

Algorithm 2 Sampling

1:
$$\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

2: for $t = T, ..., 1$ 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



Model	FID-5K \downarrow	IS ↑	$\mathbf{CA}\uparrow$	$\mathbf{PD}\downarrow$	Fool rate \uparrow
Prior Work					
pix2pix [†]	24.41	-	-	-	-
PixColor [‡]	24.32	-	-	-	29.90%
Coltran ^{††}	19.37	-	-	-	36.55%
This paper					
Regression	17.89	169.8	68.2%	60.0	39.45%
Palette	15.78	200.8	72.5%	46.2	47.80%
Original images	14.68	229.6	75.6%	0.0	-

Table 1: Colorization quantitative scores and fool rates on ImageNet val set indicate that Palette outputs are bridging

Model		Imag	Places2			
	FID ↓	IS ↑	CA ↑	PD ↓	FID ↓	PD ↓
20-30% free form						
DeepFillv2 [Yu et al. 2019]	9.4	174.6	68.8%	64.7	13.5	63.0
HiFill [Yi et al. 2020]	12.4	157.0	65.7%	86.2	15.7	92.8
Co-ModGAN [Zhao et al. 2021]	-	-	-	-	12.4	51.6
Palette (Ours)	5.2	205.5	72.3%	27.6	11.7	35.0
128×128 center						
DeepFillv2 [Yu et al. 2019]	18.0	135.3	64.3%	117.2	15.3	96.3
HiFill [Yi et al. 2020]	20.1	126.8	62.3%	129.7	16.9	115.4
Co-ModGAN [Zhao et al. 2021]	-	-	-	-	13.7	86.2
Palette (Ours)	6.6	173.9	69.3%	59.5	11.9	57.3
Original images	5.1	231.6	74.6%	0.0	11.4	0.0

Table 2: Quantitative evaluation for free-form and center inpainting on ImageNet and Places2 validation images.

Denoising Diffusion Restoration Models(DDRM)

Background:

- Previous efficient solutions often require problem-specific supervised training to model the posterior.
- DDRM is an efficient, not problem-specific ,unsupervised posterior sampling method.

Problem Define:

Many tasks in image restoration can be cast as linear inverse problems.

y = Hx + z, where *H* is a known linear degradation matrix, $z \sim N(0, \sigma^2 I)$

Main idea of DDRM:



(Independent of inverse problem)

(Dependent on inverse problem)

DDRM is to leverage the singular value decomposition of H and transform both x and the possibly noisy y to a shared spectral space.

$$egin{aligned} m{H} = m{U} m{\Sigma} m{V}^ op & ar{\mathbf{x}}_t = m{V}^ op m{\mathbf{x}}_t & ar{\mathbf{y}} = m{\Sigma}^\dagger m{U}^ op m{\mathbf{y}} \end{aligned}$$

Variational distribution

Reverse diffusion process $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_{t-1}; \tilde{\boldsymbol{\mu}}(\mathbf{x}_t, \mathbf{x}_0), \tilde{\boldsymbol{\beta}}_t \mathbf{I})$

 $\bar{\mathbf{x}}_{\theta,t}^{(i)}$ represent this prediction made by a model $f(x_{t+1}, t+1): \mathbb{R}^{n*n} \times \mathbb{R} \to \mathbb{R}^{n*n}$

JPEG Artifact Correction using Denoising Diffusion Restoration Models

Bahjat Kawar*

Technion, Haifa, Israel bahjat.kawar@cs.technion.ac.il

Stefano Ermon

Stanford University, CA ermon@cs.stanford.edu Jiaming Song* NVIDIA jiamings@nvidia.com

Michael Elad

Technion, Haifa, Israel elad@cs.technion.ac.il

Non-linear DDRM

For the case of no noise in the observation y, the general DDRM process to sample from $p_{\theta}^{(t)}(\mathbf{x}_t|\mathbf{x}_{t+1},\mathbf{y})$ for linear inverse problems simplifies to be

$$\mathbf{x}_{t}' = f_{\theta}^{(t+1)}(\mathbf{x}_{t+1}) - \boldsymbol{H}^{\dagger}\boldsymbol{H}f_{\theta}^{(t+1)}(\mathbf{x}_{t+1}) + \boldsymbol{H}^{\dagger}\mathbf{y},$$

$$\mathbf{x}_{t} = \sqrt{\alpha_{t}} \left(\eta_{b}\mathbf{x}_{t}' + (1-\eta_{b})f_{\theta}^{(t+1)}(\mathbf{x}_{t+1})\right) + \sqrt{1-\alpha_{t}} \left(\eta\epsilon_{t} + (1-\eta)\epsilon_{\theta}^{(t+1)}(\mathbf{x}_{t+1})\right)$$

if we treat H as the JPEG encoding operator, H^+ as the JPEG decode operator

With this insight, we can simply perform JPEG restoration with DDRM with the update rule predicted noise
$$\mathbf{x}'_{t} = f_{\theta}^{(t+1)}(\mathbf{x}_{t+1}) - \text{Decode}\left(\text{Encode}\left(f_{\theta}^{(t+1)}(\mathbf{x}_{t+1})\right)\right) + \text{Decode}\left(\mathbf{y}\right),$$
$$\mathbf{x}_{t} = \sqrt{\alpha_{t}}\left(\eta_{b}\mathbf{x}'_{t} + (1-\eta_{b})f_{\theta}^{(t+1)}(\mathbf{x}_{t+1})\right) + \sqrt{1-\alpha_{t}}\left(\eta_{\epsilon}t\right) + (1-\eta_{\epsilon}\epsilon_{\theta}^{(t+1)}(\mathbf{x}_{t+1})\right),$$

which can be used in realistic settings as the quantization matrices are stored within the JPEG files.

NLO.1)

























