

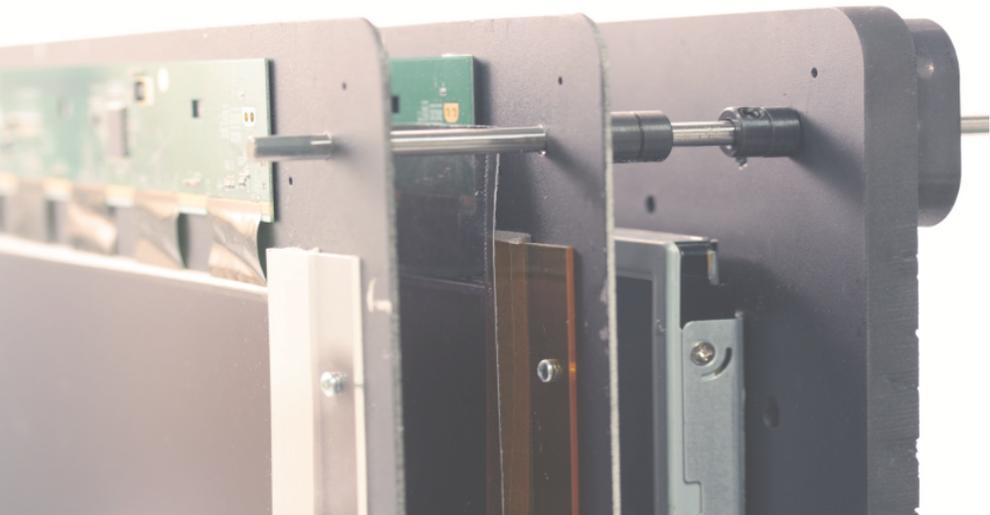
Solving Inverse Problems with Diffusion Model-based Priors

EE367/CS448I: Computational Imaging

stanford.edu/class/ee367

Lecture 13

Gordon Wetzstein
Stanford University



Recap

- Model data distribution using score function (gradient of log prior)
- Score function can be learned robustly via denoising score matching
- Diffusion models can be represented as SDEs/ODEs → rich literature on solvers, sampling, etc.
- Generate new examples using sampling

Overview

- Diffusion Model Sampling (continued)
- Image Editing/Translation with SDEdit
- Solving Inverse Problems with Posterior Sampling

Diffusion Model Sampling

(continued)

(Continuous) Diffusion Models as SDEs

- Forward SDE: $d\mathbf{x} = \mathbf{f}(\mathbf{x}, t)dt + g(t)d\mathbf{w}$
 - drift coefficient \nearrow (pointing to $\mathbf{f}(\mathbf{x}, t)$)
 - diffusion coefficient \nwarrow (pointing to $g(t)$)
 - noise with variance dt \swarrow (pointing to $d\mathbf{w}$)
- Reverse SDE: $d\mathbf{x} = [\mathbf{f}(\mathbf{x}, t) - g^2(t)\nabla_x \log p_t(\mathbf{x})]dt + g(t)d\tilde{\mathbf{w}}$
 - standard Wiener process, i.e., noise with variance dt \nearrow (pointing to $d\tilde{\mathbf{w}}$)

(Continuous) Diffusion Models as SDEs

- Forward SDE: $d\mathbf{x} = \mathbf{f}(\mathbf{x}, t)dt + g(t)d\mathbf{w}$
- Discretize SDE using the Euler-Maruyama (or Euler) method:

$$x_t = x_{t-1} + \mathbf{f}(x_{t-1}, t)\Delta t + g(t)\sqrt{\Delta t}\epsilon$$

- Same idea for reverse SDE

↑
noise, $\epsilon \sim \mathcal{N}(0,1)$

Denoising Diffusion Probabilistic Models (DDPM)

- DDPM is a specific form of the general, discretized SDE
- Variance-preserving forward diffusion: $x_t = \sqrt{1 - \beta_t} x_{t-1} + \sqrt{\beta_t} \epsilon_{t-1}$



noise schedule, i.e., how much noise to add at time step t

Denoising Diffusion Probabilistic Models (DDPM)

- DDPM is a specific form of the general, discretized SDE
- Variance-preserving forward diffusion: $x_t = \sqrt{1 - \beta_t} x_{t-1} + \sqrt{\beta_t} \epsilon_{t-1}$
- One-step forward diffusion:

$$x_t = \sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon \sim \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t} x_0, (1 - \bar{\alpha}_t)I)$$

$$\alpha_t = 1 - \beta_t, \quad \bar{\alpha}_t = \prod_{s=1}^t \alpha_s, \quad \epsilon \sim \mathcal{N}(0, I)$$

- Reverse diffusion: $x_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_{\theta}(x_t, t) \right) + \sqrt{\beta_t} \epsilon_t$

Denoising Diffusion Probabilistic Models (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
$$\nabla_{\theta} \|\boldsymbol{\epsilon} - \boldsymbol{\epsilon}_{\theta}(\sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \boldsymbol{\epsilon}, t)\|^2$$
 - 6: **until** converged
-

train a noise estimation network!

Algorithm 2 Sampling

- 1: $\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
 - 2: **for** $t = T, \dots, 1$ **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
-

Denoising Diffusion Probabilistic Models (DDPM)

- DDPM is a specific form of the general SDE :

$$\mathbf{f}(\mathbf{x}, t) = -\frac{1}{2}\beta(t)\mathbf{x}, \quad g(t) = \sqrt{\beta(t)}$$

- Not straightforward to see relationship between SDE and update rules via Euler; see Song et al., Appendix B (Eqs. 22-25) for detailed derivation
- The DDPM sampling scheme requires many steps (e.g., $T=1,000$ or more)

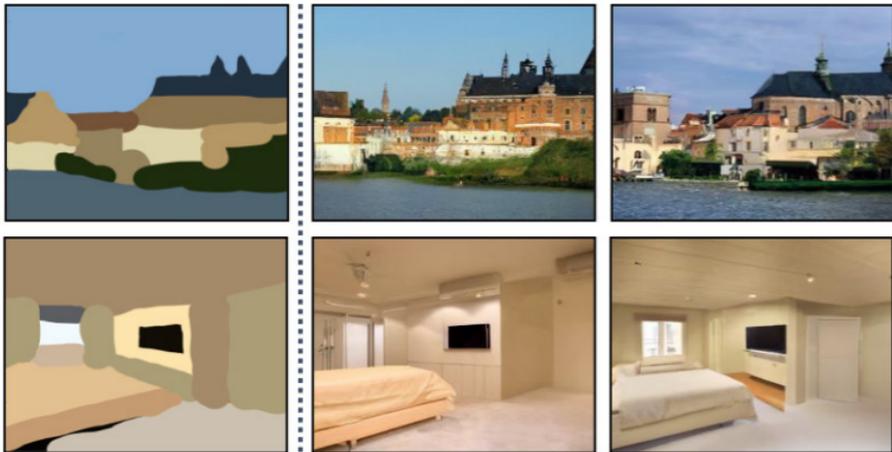
Other SDE Formulations / Sampling Schemes

- DDIM (Song et al.) significantly reduces the number of steps compared with DDPM using a non-Markovian derivation
- Karras et al. derived a unified SDE-based formulation that leads to a family of viable drift & diffusion coefficients, with both DDPM & DDIM being special cases
- Rich literature on sampling schemes; e.g., StableDiffusion (<SD3) used DPM-Solver++ (Lu et al.)
- Newer versions of diffusion models, including FLUX and SD3, use a slightly different mechanism - flow matching (Lipman et al.; Hawley)

Image Editing/Translation via SDEdit

SDEdit

Stroke Painting to Image



Input

Output

Image Compositing



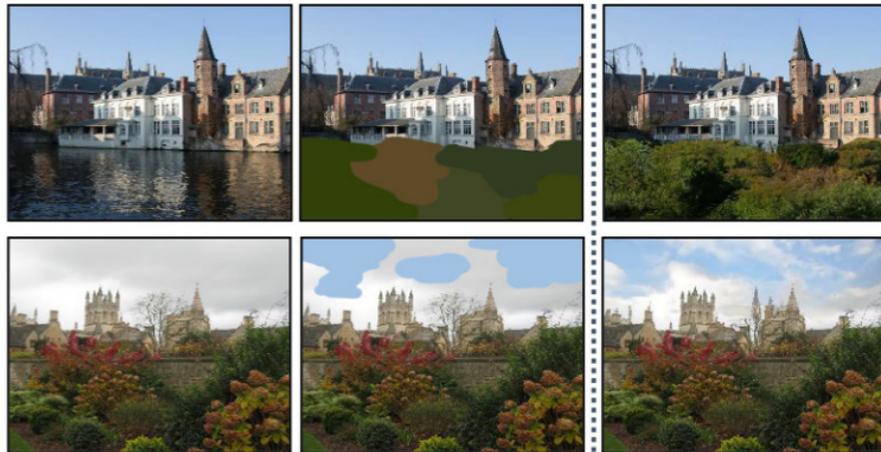
Source

Output

Source

Output

Stroke-based Editing



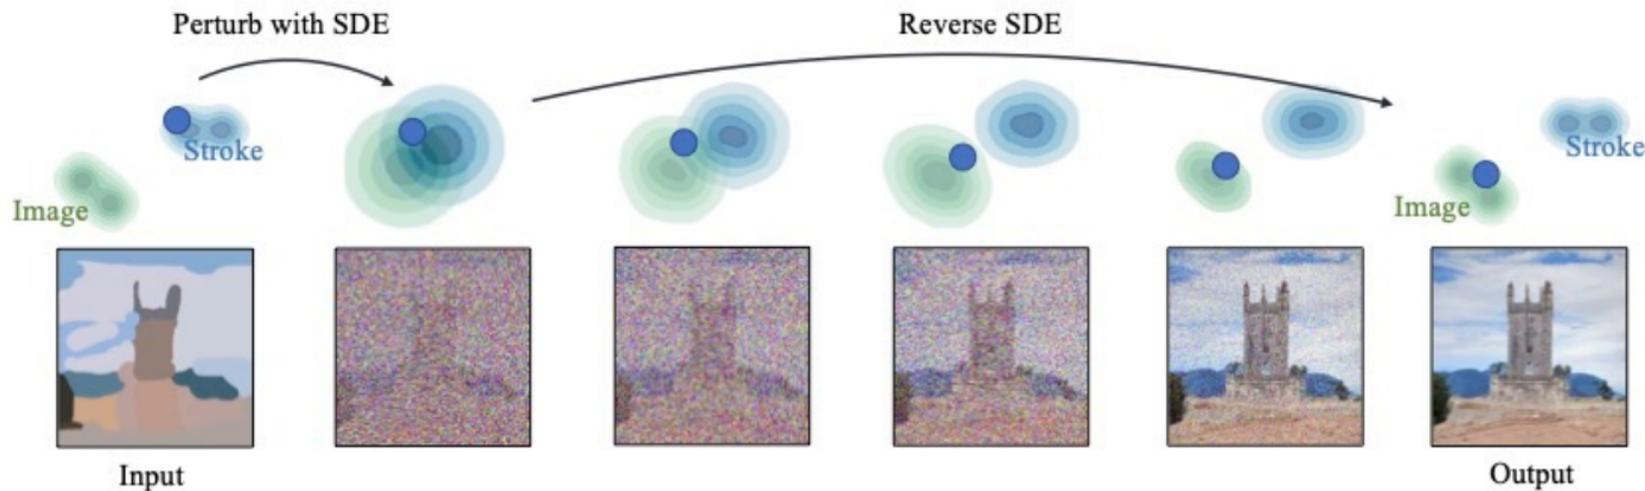
Source

Input

Output

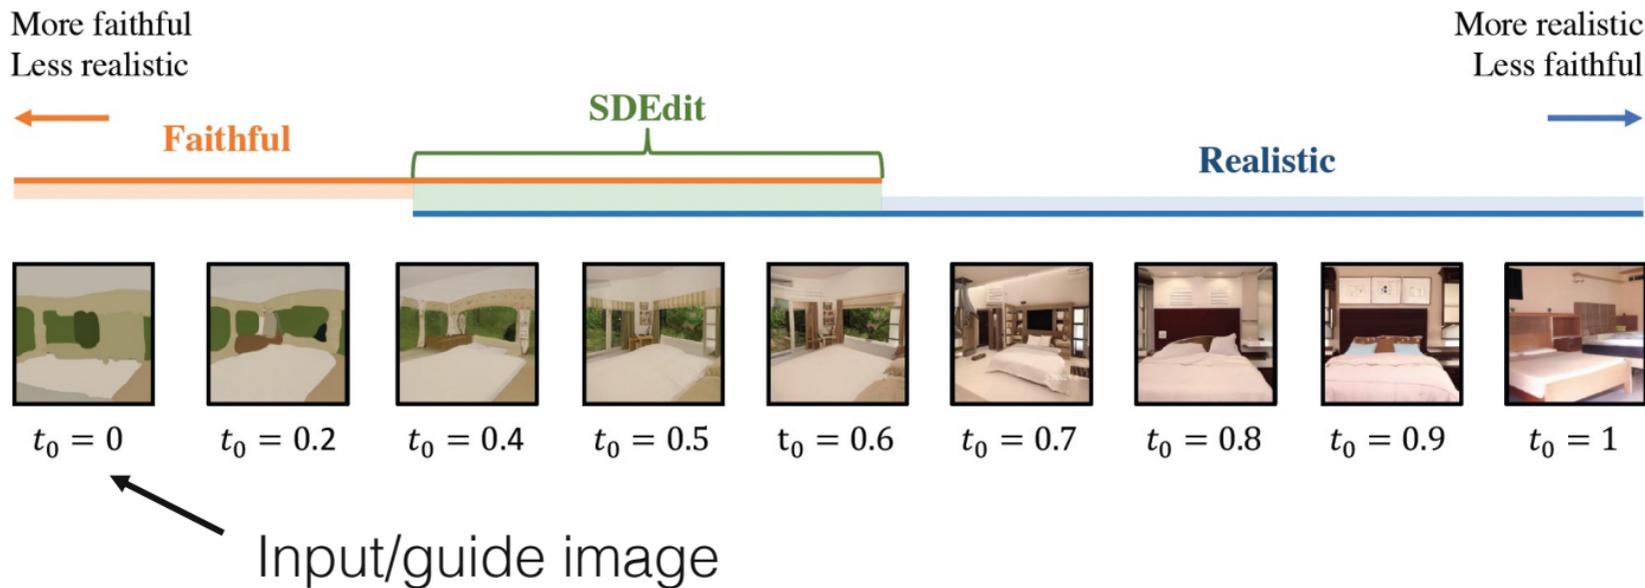
SDEdit

1. Add noise to input image via forward diffusion process
2. Apply reverse diffusion process



SDEdit

Question: how much noise to add?



SDEdit

- SDEdit doesn't solve inverse problems, but it's a simple-enough image editing, or image-to-image translation, approach that it might be useful
- To solve an inverse problem, we need to take the image formation model into account!

A Bayesian Perspective of Inverse Problems

A Bayesian Perspective of Inverse Problems

...from lecture 10...

- Image formation model: $\mathbf{b} = \mathbf{A}\mathbf{x} + \boldsymbol{\eta}$, $\mathbf{b} \in \mathbb{R}^M$, $\mathbf{x} \in \mathbb{R}^N$, $\mathbf{A} \in \mathbb{R}^{M \times N}$

- Interpret as random variables:
 $\mathbf{x}_i \sim \mathcal{N}(\mathbf{x}_i, 0)$, $\boldsymbol{\eta}_i \sim \mathcal{N}(0, \sigma^2)$
 $\mathbf{b}_i \sim \mathcal{N}((\mathbf{A}\mathbf{x})_i, \sigma^2)$

- Probability of observation i :
$$p(\mathbf{b}_i | \mathbf{x}_i, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(\mathbf{b}_i - (\mathbf{A}\mathbf{x})_i)^2}{2\sigma^2}}$$

- Joint probability of all observations:
$$p(\mathbf{b} | \mathbf{x}, \sigma) = \prod_{i=1}^M p(\mathbf{b}_i | \mathbf{x}_i, \sigma) \propto e^{-\frac{\|\mathbf{b} - \mathbf{A}\mathbf{x}\|_2^2}{2\sigma^2}}$$

ML

- simple to compute, e.g., using closed-form solution or (stochastic) gradient descent
- considers physical image formation model
- doesn't consider priors
- Only suitable for well-posed, but not ill-posed inverse problems

MAP

- Solve using plug and play, e.g., using ADMM (see lectures 10, 11)
- considers physical image formation model
- considers priors
- “point estimate”, i.e., “gives a single solution”
- best for “slightly to moderately” ill-posed inverse problems

Posterior Sampling

- Solve e.g. using guided diffusion model sampling (this lecture)
- considers physical image formation model
- considers priors
- gives a different feasible solution each time
- best for “moderately to highly” ill-posed inverse problems

Posterior Sampling with Diffusion Models

Naive Approach

- Train a new diffusion model on $p_{data}(\mathbf{x}|\mathbf{b})$ instead of $p_{data}(\mathbf{x})$?
- This is usually not feasible, because $p_{data}(\mathbf{x}|\mathbf{b})$ depends on measurement operator \mathbf{A} and the diffusion model would have to be trained for each operator \mathbf{A} !
- Instead: work with diffusion model pretrained on $p_{data}(\mathbf{x})$

Sampling with Guidance

- Reverse SDE: $d\mathbf{x} = [\mathbf{f}(\mathbf{x}, t) - g^2(t)\nabla_x \log p_t(\mathbf{x})]dt + g(t)d\tilde{\mathbf{w}}$
- Replace unconditional distribution with conditional using Bayes' rule: $p_t(\mathbf{x}|\mathbf{b}) \propto p_t(\mathbf{b}|\mathbf{x})p_t(\mathbf{x})$
- Intuition: “guide” the diffusion model sampling by the measurements \mathbf{b} to sample from $p_{data}(\mathbf{x}|\mathbf{b})$ instead of $p_{data}(\mathbf{x})$

Sampling with Guidance

- Reverse SDE: $d\mathbf{x} = [\mathbf{f}(\mathbf{x}, t) - g^2(t)\nabla_x \log p_t(\mathbf{x})]dt + g(t)d\tilde{\mathbf{w}}$
- Replace unconditional distribution with conditional using

Bayes' rule: $p_t(\mathbf{x}|\mathbf{b}) \propto p_t(\mathbf{b}|\mathbf{x})p_t(\mathbf{x})$

- Conditional reverse SDE:

$$d\mathbf{x} = [\mathbf{f}(\mathbf{x}, t) - g^2(t)(\nabla_x \log p_t(\mathbf{x}) + \nabla_x \log p_t(\mathbf{b}|\mathbf{x}))]dt + g(t)d\tilde{\mathbf{w}}$$

Sampling with Guidance

- Likelihood of measurements: $p(\mathbf{b}|\mathbf{x}, \sigma) \propto e^{-\frac{\|\mathbf{b}-\mathbf{Ax}\|_2^2}{2\sigma^2}}$
- Gradient of log-likelihood:

distribution of
"clean" images

$$\nabla_x \log p(\mathbf{b}|\mathbf{x}) = -\nabla_x \frac{\|\mathbf{b} - \mathbf{Ax}\|_2^2}{2\sigma^2} = -\frac{1}{\sigma^2} (\mathbf{A}^T (\mathbf{b} - \mathbf{Ax}))$$

- Problem: $\nabla_x \log p(\mathbf{b}|\mathbf{x}_0) \neq \nabla_x \log p_t(\mathbf{b}|\mathbf{x}_t)$

$$d\mathbf{x} = [\mathbf{f}(\mathbf{x}, t) - g^2(t)(\nabla_x \log p_t(\mathbf{x}) + \nabla_x \log p_t(\mathbf{b}|\mathbf{x}))]dt + g(t)d\tilde{\mathbf{w}}$$

distribution of "noisy" images

Sampling with Guidance

- Problem: $\nabla_x \log p(\mathbf{b}|\mathbf{x}_0) \neq \nabla_x \log p_t(\mathbf{b}|\mathbf{x}_t)$
- Problem: likelihood of the measurements is given by intractable integral: $p_t(\mathbf{b}|\mathbf{x}_t) = \int p(\mathbf{b}|\mathbf{x}_0)p(\mathbf{x}_0|\mathbf{x}_t)d\mathbf{x}_0$
- Solution: need to approximate $\nabla_x \log p_t(\mathbf{b}|\mathbf{x}_t)$ somehow!

Score ALD

- Problem: $\nabla_x \log p(\mathbf{b}|\mathbf{x}_0) \neq \nabla_x \log p_t(\mathbf{b}|\mathbf{x}_t)$
- Approach: $\nabla_x \log p(\mathbf{b}|\mathbf{x}_0) \approx \nabla_x \log p_t(\mathbf{b}|\mathbf{x}_t)$

$$\approx -\frac{1}{\sigma^2 + \gamma_t^2} (\mathbf{A}^T (\mathbf{b} - \mathbf{A}\mathbf{x}))$$



guidance strength - hyperparameter

Diffusion Posterior Sampling (DPS)

- Problem: $\nabla_x \log p(\mathbf{b}|\mathbf{x}_0) \neq \nabla_x \log p_t(\mathbf{b}|\mathbf{x}_t)$
- Approach: $\nabla_x \log p(\mathbf{b}|\mathbf{x}_0) \approx \nabla_x \log p_t(\mathbf{b}|\mathbf{x}_0 = \mathbb{E}[\mathbf{x}_0|\mathbf{x}_t])$
- Tweedie's formula for variance preserving (i.e., DDPM) discrete forward diffusion (see Eq. 10 in Chung et al.):

$$\hat{x}_0 = \mathbb{E}[\mathbf{x}_0|\mathbf{x}_t] = \frac{1}{\sqrt{\bar{\alpha}_t}} (x_t + (1 - \bar{\alpha}_t) \nabla_{x_t} \log p_t(x_t))$$

DDPM vs. DPS

DDPM: noise prediction vs. score prediction

DDPM reverse diffusion w/ noise predictor

$\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
for $t = T, \dots, 1$ **do**
 $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $t > 1$, else $\mathbf{z} = \mathbf{0}$
 $x_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(x_t + \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_{\theta}(x_t, t) \right) + \sqrt{1 - \alpha_t} \mathbf{z}$
end for
return \mathbf{x}_0

where $\bar{\alpha}_t \triangleq \prod_{i=1}^t \alpha_i$ and $\alpha_t \triangleq 1 - \beta_t$

DDPM reverse diffusion w/ score predictor

$\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
for $t = T, \dots, 1$ **do**
 $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $t > 1$, else $\mathbf{z} = \mathbf{0}$
 $x_{t-1} = \frac{1}{\sqrt{\alpha_t}} (x_t + (1 - \alpha_t) \mathbf{s}_{\theta}(x_t, t)) + \sqrt{1 - \alpha_t} \mathbf{z}$
end for
return \mathbf{x}_0

- Noise prediction formulation is used in original paper; score prediction formulation in project
- Both formulations are equivalent

DDPM: score prediction variants

DDPM reverse diffusion w/ score predictor

$\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
for $t = T, \dots, 1$ **do**
 $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $t > 1$, else $\mathbf{z} = \mathbf{0}$
 $x_{t-1} = \frac{1}{\sqrt{\alpha_t}}(x_t + (1 - \alpha_t)\mathbf{s}_\theta(x_t, t)) + \sqrt{1 - \alpha_t}\mathbf{z}$
end for
return \mathbf{x}_0

where $\bar{\alpha}_t \triangleq \prod_{i=1}^t \alpha_i$ and $\alpha_t \triangleq 1 - \beta_t$

DDPM reverse diffusion w/ score predictor

$\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
for $t = T, \dots, 1$ **do**
 $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $t > 1$, else $\mathbf{z} = \mathbf{0}$
 $\hat{\mathbf{x}}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}}(x_t + (1 - \bar{\alpha}_t)\mathbf{s}_\theta(x_t, t))$
 $x_{t-1} = \frac{\sqrt{\alpha_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t}x_t + \frac{\sqrt{\bar{\alpha}_{t-1}}(1 - \alpha_t)}{1 - \bar{\alpha}_t}\hat{\mathbf{x}}_0 + \sqrt{1 - \alpha_t}\mathbf{z}$
end for
return \mathbf{x}_0

- Variant on right is slightly more complicated, but directly shows how the predicted clean image $\hat{\mathbf{x}}_0$ is part of this

DDPM vs. DPS

DDPM reverse diffusion w/ score predictor

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

for $t = T, \dots, 1$ **do**

$$\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \text{ if } t > 1, \text{ else } \mathbf{z} = \mathbf{0}$$

$$\hat{\mathbf{x}}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}} (x_t + (1 - \bar{\alpha}_t) \mathbf{s}_\theta(x_t, t))$$

$$x_{t-1} = \frac{\sqrt{\alpha_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} x_t + \frac{\sqrt{\bar{\alpha}_{t-1}}(1 - \alpha_t)}{1 - \bar{\alpha}_t} \hat{\mathbf{x}}_0 + \sqrt{1 - \alpha_t} \mathbf{z}$$

end for

return \mathbf{x}_0

DPS reverse diffusion

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

for $t = T, \dots, 1$ **do**

$$\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \text{ if } t > 1, \text{ else } \mathbf{z} = \mathbf{0}$$

$$\hat{\mathbf{x}}_0 = \frac{1}{\sqrt{\bar{\alpha}_t}} (x_t + (1 - \bar{\alpha}_t) \mathbf{s}_\theta(x_t, t))$$

$$x'_{t-1} = \frac{\sqrt{\alpha_t}(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} x_t + \frac{\sqrt{\bar{\alpha}_{t-1}}(1 - \alpha_t)}{1 - \bar{\alpha}_t} \hat{\mathbf{x}}_0 + \sqrt{1 - \alpha_t} \mathbf{z}$$

$$x_{t-1} = x'_{t-1} - \zeta_t \nabla_{x_t} \|\mathbf{b} - \mathbf{A} \hat{\mathbf{x}}_0\|_2^2$$

end for

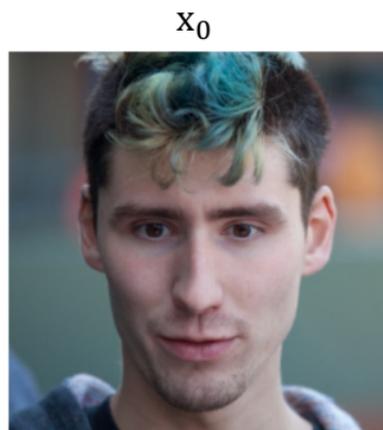
return \mathbf{x}_0

- DPS is equivalent to DDPM with one additional step!

Results from Diffusion Project

Diffusion Project

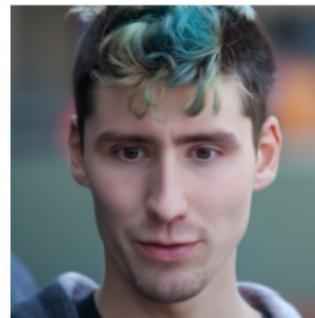
Task 2: Single-step Denoising



t=30



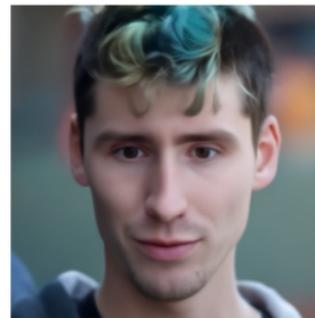
Denoised



PSNR/LPIPS

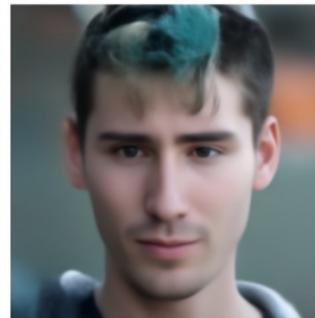
37.4
0.0378

t=100



32.7
0.0941

t=300



27.3
0.203

Diffusion Project

Task 3: Unconditional Image Generation

Example results



Diffusion Project

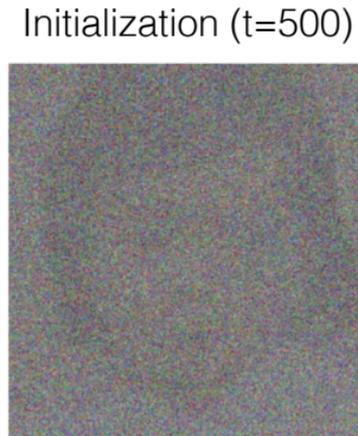
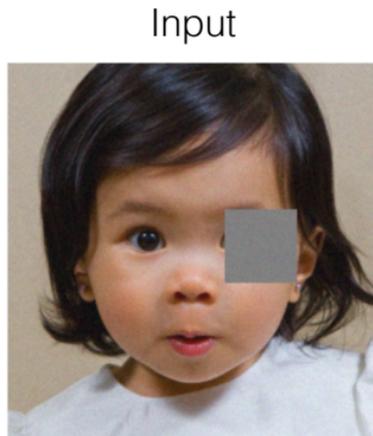
Task 4: SDEdit for Inverse Problems

Inpainting

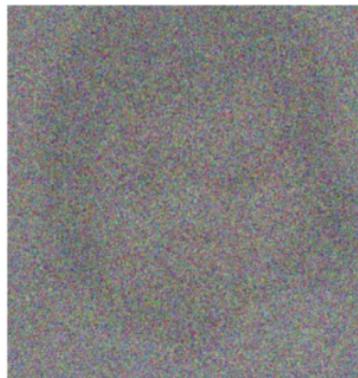
Ground Truth



Deconvolution



20.2 / 0.189



20.1 / 0.232

Diffusion Project

Task 5: ScoreALD for Inverse Problems

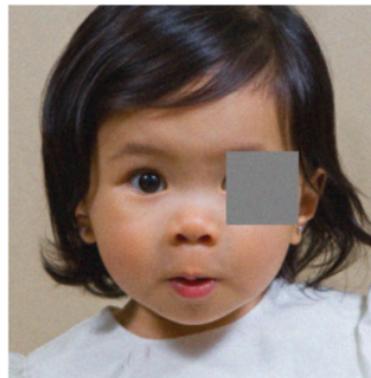
PSNR/LPIPS

Input

Reconstruction

Inpainting

Ground Truth



28.5 / 0.0605

Deconvolution



22.2 / 0.144

Diffusion Project

Task 6: DPS for Inverse Problems

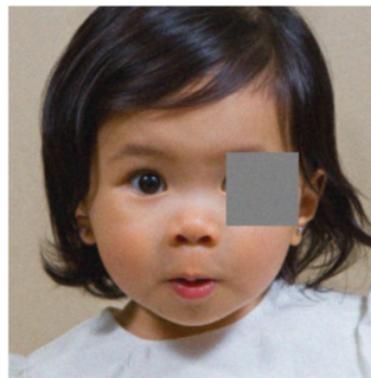
PSNR/LPIPS

Ground Truth



Inpainting

Input



Reconstruction



34.0 / 0.0239

Deconvolution



28.4 / 0.0518

Summary

- DPS works for linear and nonlinear inverse problems with Gaussian or Poisson noise! Also easy to understand and implement. That's why we discussed it
- However, Score-ALD and DPS are two of *many* approaches to approximate $\nabla_x \log p_t(\mathbf{b}|\mathbf{x}_t)$
- See Daras et al. "A Survey on Diffusion Models for Inverse Problems", arxiv 2024, for a great and comprehensive survey

References and Further Reading

Diffusion Models:

- Y. Song, S. Ermon, "Generative modeling by estimating gradients of the data distribution", NeurIPS 2019
- **J. Ho, A. Jain, P. Abbeel, "Denoising Diffusion Probabilistic Models", NeurIPS 2020 (DDPM)**
- J. Song, C. Meng, S. Ermon, "Denoising diffusion implicit models", ICLR, 2021 (DDIM)
- Y. Song, J. Sohl-Dickstein, D. Kingma, A. Kumar, S. Ermon, B. Poole, "Score-based generative modeling through stochastic differential equations", ICLR 2021
- T. Karras, M. Aittala, T. Aila, S. Laine, "Elucidating the Design Space of Diffusion-Based Generative Models", NeurIPS 2022
- **C. Meng, Y. He, Y. Song, J. Song, J. Wu, J.Y. Zhu, S. Ermon, "SDEdit: Guided Image Synthesis and Editing with Stochastic Differential Equations", ICLR 2022**
- C. Lu, Y. Zhou, F. Bao, J. Chen, C. Li, J. Zhu, "DPM-Solver: A Fast ODE Solver for Diffusion Probabilistic Model Sampling in Around 10 Steps", NeurIPS 2022
- C. Lu, Y. Zhou, F. Bao, J. Chen, C. Li, J. Zhu, "DPM-Solver++: Fast Solver for Guided Sampling of Diffusion Probabilistic Models", arxiv 2022
- Y. Lipman, R. Chen, H. Ben-Hamu, M. Nickel, M. Le, "Flow Matching for Generative Modeling", ICLR 2023
- **G. Daras, H. Chung, C. Lai, Y. Mitsufuji, J. Ye, P. Milanfar, A. Dimakis, M. Delbracio, "A Survey on Diffusion Models for Inverse Problems", arxiv 2024**
- **A. Jalal, M. Arvinte, G. Daras, E. Price, A. Dimakis, J. Tamir, "Robust Compressed Sensing MRI with Deep Generative Priors", NeurIPS 2021**
- H. Chung, B. Sim, D. Ryu, J. Ye, "Improving Diffusion Models for Inverse Problems using Manifold Constraints", NeurIPS 2022
- **H. Chung, J. Kim, M. McCann, M. Klasky, J. Ye, "Diffusion Posterior Sampling for General Noisy Inverse Problems", ICLR 2023**
- **S. Hawley, "Flow with what you know", <https://drscotthawley.github.io/blog/posts/FlowModels.html>, blog post 2024**

Plug & play image restoration – original paper and extension for diffusion models

- S. Venkatakrishnan, C. Bouman, B. Wohlberg. Plug-and-play priors for model based reconstruction. In 2013 IEEE Global Conference on Signal and Information Processing, 2013
- Y. Zhu, K. Zhang, J. Liang, J. Cao, B. Wen, R. Timofte, L. Gool. Denoising diffusion models for plug-and-play image restoration, CVPR 2023