Diffusion Methods for Solving Inverse Problems in Computational Imaging

Karthik Karumanchi

Motivation

Solving inverse problems (in-painting, deconvolution) in computational imaging in the presence of noise is ill-conditioned

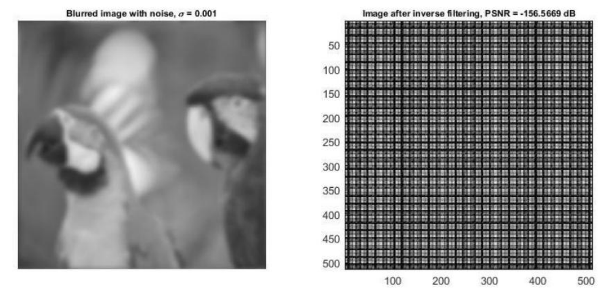

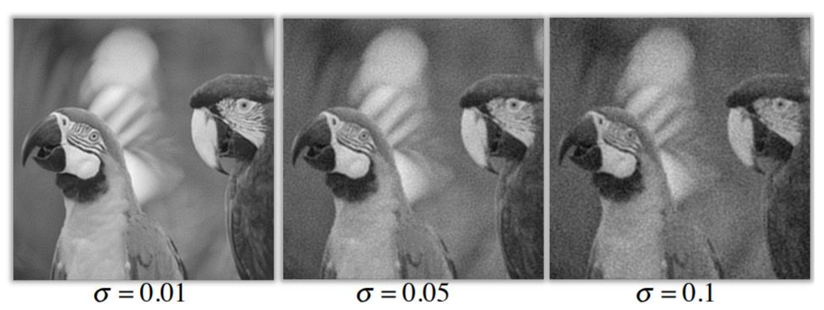

Conventional methods such as Wiener filtering or TV priors do not work well with complex images

Diffusion-based methods aim to learn a noise to manifold map and improve reconstruction quality

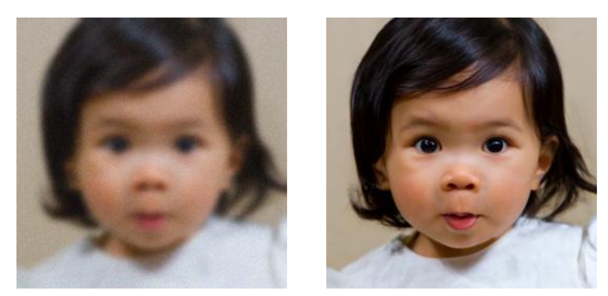

We present and compare three algorithms using the same diffusion network to deblur and in-paint images

Diffusion-based algorithms: SDEdit, ScoreALD and DPS

SDEdit

Small amount of noise is added to the input image and gradually denoised through the reverse diffusion process

Pros:
- Computationally inexpensive

Cons:
- Ignores the image formation model
- May hallucinate and generate dissimilar images

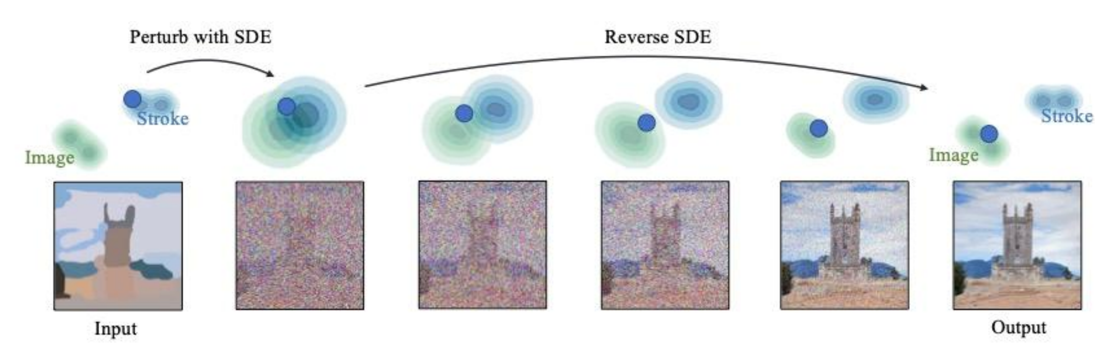

ScoreALD

Treats inverse problems as a posterior sampling task by adding a gradient descent step on the measurement . Utilizes an annealing factor to balance between the generative prior and data consistency.

Pros:
- Provides a mathematically grounded way to incorporate measurements.

Cons:
- Highly sensitive to annealing factor schedule

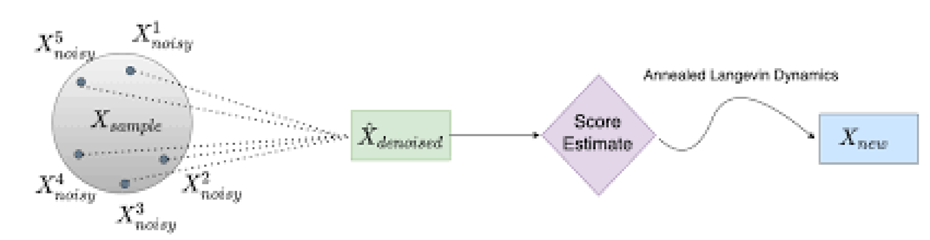

DPS

Computes gradients with respect to the estimated clean image. This allows the model to apply the degradation operator to a manifold-constrained estimate rather than a noisy intermediate, leading to higher stability.

Pros:
- Robust for general noisy inverse problems

Cons:
- Computationally expensive due to the need for backpropagation at every timestep
- Requires a carefully tuned scale parameter

Related Work

Algorithm	Features
Wiener Filtering	Does not require model training, but does not work on higher noise levels
TV Priors	Needs a handcrafted prior and cannot preserve high frequency variation
SDEdit[1]	Can handle increased noise levels but may not retain features of the original image
Posterior Sampling[2],[3]	Combines features of diffusion with the image formation model to solve inverse problems

Experimental Results

Input | SDEdit | ScoreALD | PSNR

In-painting

Deblurring

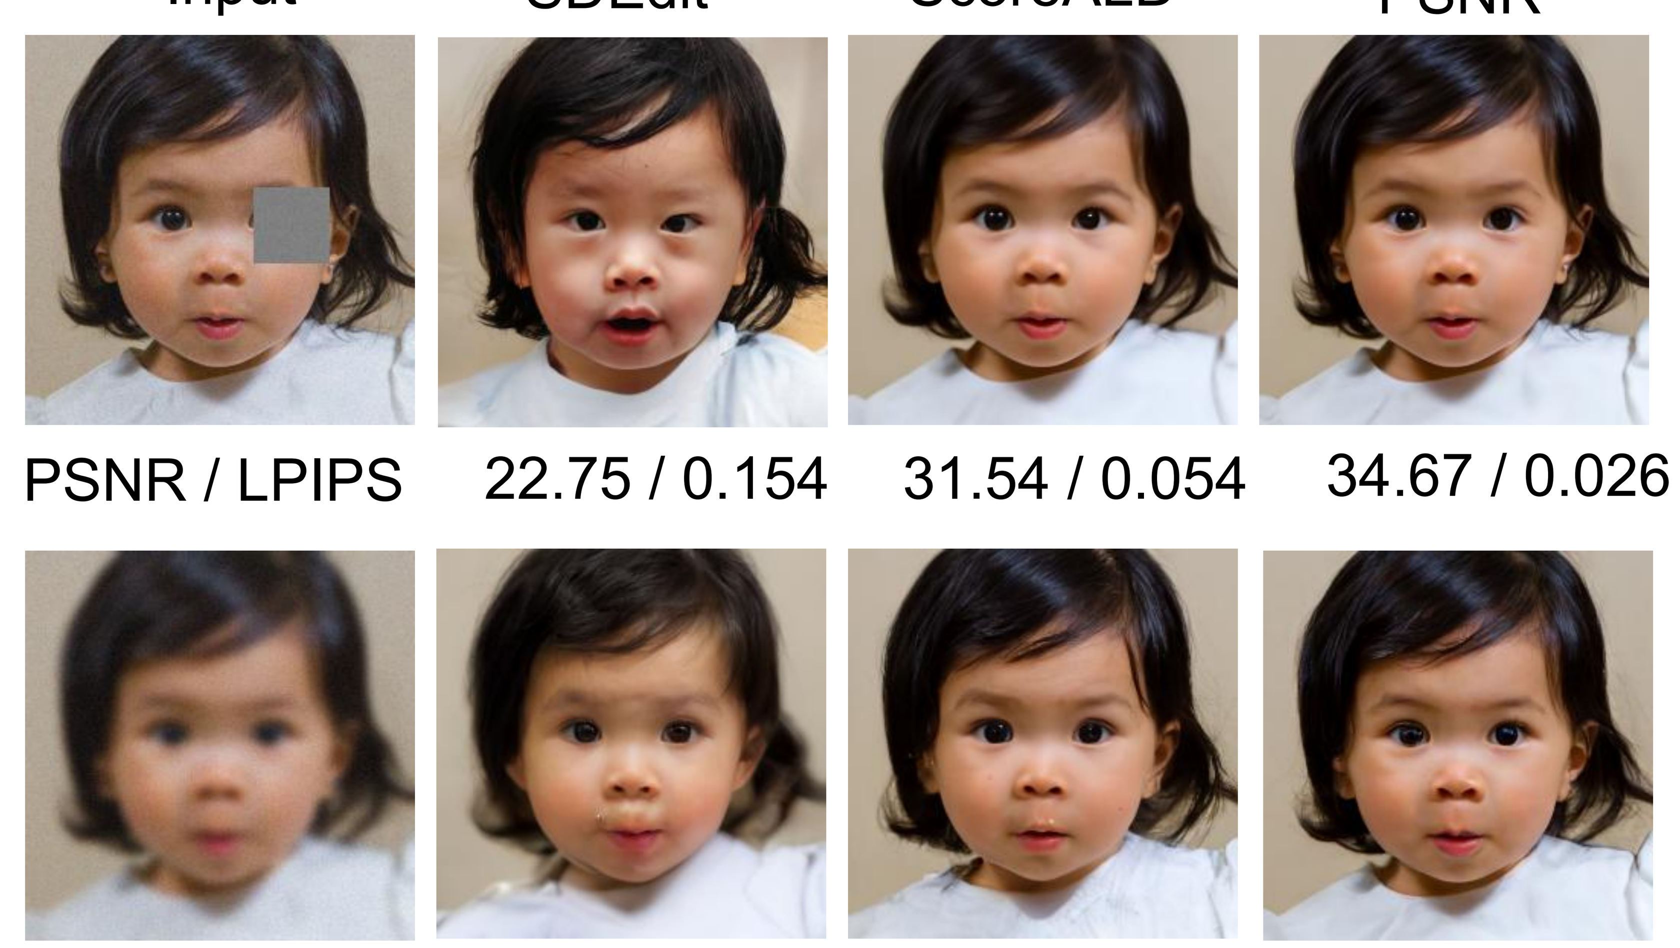

	Input	SDEdit	ScoreALD	PSNR
In-painting	PSNR / LPIPS	22.75 / 0.154	31.54 / 0.054	34.67 / 0.026
Deblurring	PSNR / LPIPS	23.58 / 0.17	25.82 / 0.114	28.50 / 0.056

- DPS performs better than SDEdit and ScoreALD across in-painting and deblurring tasks
- Since SDEdit does not account for the image formation model, its outputs look visually different from the input
- Images on the left are the best results for each algorithm across different hyperparameters. A suboptimal hyperparameter leads to significant degradation of performance.

References

[1] C. Meng et al. ,"Sdedit: Guided image synthesis and editing with stochastic differential equations" ICLR, 2022
[2] Jalal, M. Arvinte, G. Daras, E. Price, A. Dimakis, and J. Tamir, "Robust compressed sensing MRI with deep generative priors,"NeurIPS, 2021
[3] H. Chung et al. "Diffusion posterior sampling for general noisy inverse problems," ICLR, 2023