ARIA
APPLIED RESEARCH IN ACTION

Unsupervised Real Image Denoising

Using a pre-trained off-the-shelf diffusion model as an image prior for unsupervised real-world image denoising.

Tianshu Kuai

Igor Gilitschenski
ACADEMIC SUPERVISOR

Alex Levinshtein
INDUSTRY SUPERVISOR

Noisy Image | Denoised Image | Enhanced Image

PROJECT SUMMARY

Image denoising has long been an important task in computer vision. Recent deep learning based methods have shown remarkable performance, particularly when trained on large-scale paired datasets in a supervised fashion. However, models trained on synthetic data often fail on real-world noisy images, and acquiring paired datasets for such scenarios requires significant effort. In this work, we utilize a pre-trained diffusion model as a generative prior for unsupervised real-world image denoising. Unlike many recent approaches that employ diffusion models at test time, we only do so at training time and thus maintain efficient inference time performance. Specifically, we start with a pre-trained image denoiser that is trained with paired synthetic data in a supervised fashion. The denoiser's results on real noisy samples are then enhanced by a pre-trained diffusion model at relatively small timesteps to obtain clean pseudo targets, which can be used to finetune the pre-trained denoiser. The finetuning helps the denoiser adapt to real-world noisy images. Extensive experiments show that the proposed approach can consistently improve a pre-trained denoiser's performance generalizing across several denoiser architectures. It achieves competitive performance compared to other state-of-the-art unsupervised methods when applied to denoisers that require similar or smaller computational costs.

Samsung AI Center-Toronto

Computer Science
UNIVERSITY OF TORONTO
Master of Science in Applied Computing