Multiresolution Hash Encoding as Image Prior

Tianxiang Dai txdai@stanford.edu
Department of Electrical Engineering, Stanford University

Motivation

Traditional denoising methods often rely on handcrafted priors or assumptions about the underlying image structure.

Instant Neural Graphics Primitives (instant-NGP) shows promising results in representing and reconstructing images.

These functions inherently possess certain priors that can potentially be exploited for image denoising.

The Deep Image Prior (DIP) approach has demonstrated the effectiveness of using convolutional neural networks (CNNs) as a prior for image restoration tasks.

Dataset

The Berkeley Segmentation Dataset and Benchmark (BSDS-300) is a diverse collection of 300 natural images that has been widely adopted for evaluating image denoising algorithms.

The dataset's diverse content, including varying textures, structures, and lighting conditions, presents a range of challenges for denoising algorithms, making it a valuable resource for the development and comparison of image denoising methods.

Speed Comparison

	Deep Image Prior	Instant-NGP	Frequency & MLP	MLP
ms/it	49.12	21.23	7.52	7.47

Models

Deep image prior model uses convolutional encoder-decoder with skip connections.

Instant-NGP encodes image coordinates (x,y) to hidden representation by looking up and interpolating from a hash table, and then use a MLP to decode this to RGB values.

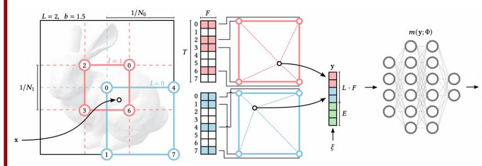

Experiments

15 images are selected from the BSDS-300 dataset. Three levels of gaussian noise are added and removed using the baseline model of deep image prior, instant-NGP model using MLP with hash table encoding, MLP with frequency encoding and pure MLP. For comparison, along with the denoised result, the moving average of the denoised result over the iterations is also measured with beta=0.99. The metric is peak signal-to-noise ratio:

$$PSNR = 10 \cdot \log_{10}\left(\frac{MAX_I^2}{MSE}\right)$$

noise level (sigma)	Deep Image Prior	Instant-NGP	Frequency & MLP	MLP
15	30.79 (it 3000)	29.59 (it 1000)	24.93 (it 1000)	11.55 (it 1000)
25	29.30 (it 3000)	26.82 (it 580)	24.62 (it 1000)	11.53 (it 1000)
50	25.52 (it 1400)	22.48 (it 380)	23.30 (it 1000)	11.54 (it 1000)

Experiment Results

Deep image prior delivers better results than instat-NGP. However, running instant-NGP to the highest PSNR only requires 1/5 the number of iterations for deep image prior, taking roughly 1/30 the time, proving it to be a more efficient denoising model than deep image prior, and this ability is closely linked to its novel hash table encoding.

Instant-NGP shows different prior than deep image prior, with more focus on details (grainy).

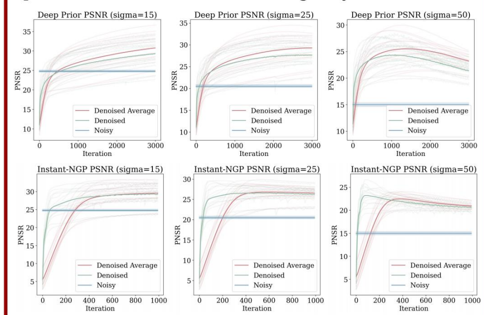

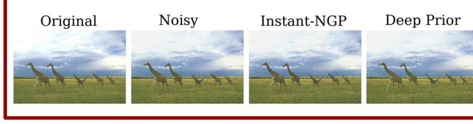

References

[1] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. International Journal of Computer Vision, 128(7):1867–1888, March 2020
[2] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. CoRR, abs/2201.05989, 2022.
[3] Martin, D., Fowlkes, C., Tal, D., & Malik, J. (2001). A Database of Human Segmented Natural Images and its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics. Proc. 8th Int'l Conf. Computer Vision, 2, 416–423.