
JPEG Artifact Correction using Denoising Diffusion Restoration Models

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

Jiaming Song*
NVIDIA
jiamings@nvidia.com

Stefano Ermon
Stanford University, CA
ermon@cs.stanford.edu

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

Abstract

Diffusion models can be used as learned priors for solving various inverse problems. However, most existing approaches are restricted to linear inverse problems, limiting their applicability to more general cases. In this paper, we build upon Denoising Diffusion Restoration Models (DDRM) and propose a method for solving some non-linear inverse problems. We leverage the pseudo-inverse operator used in DDRM and generalize this concept for other measurement operators, which allows us to use pre-trained unconditional diffusion models for applications such as JPEG artifact correction. We empirically demonstrate the effectiveness of our approach across various quality factors, attaining performance levels that are on par with state-of-the-art methods trained specifically for the JPEG restoration task. Our code is available at <https://github.com/bahjat-kawar/ddrm-jpeg>.



Figure 1: Pairs of JPEG images and restorations using our method. Best viewed zoomed in.

1 Introduction

Many image processing problems are instances of inverse problems [22, 20, 21]. In real-world applications, one would often need to face multiple different degradation models [33, 14, 16], where training problem-specific models on each case could be expensive [27]. Therefore, it is valuable to develop methods that apply problem-agnostic models, which would adapt to different degradation models at inference time without retraining. Existing approaches, while achieving decent performance on a variety of tasks, are generally limited to linear inverse problems [5, 34, 18, 16], leaving out certain important non-linear inverse problems such as JPEG artifact correction. Since JPEG is a lossy image compression format [37], JPEG images exhibit loss of quality and undesired artifacts. Several methods have been developed for addressing this problem [40, 4, 9, 23, 10, 15].

*Equal contribution.

Figure 2: Pairs of quantized (9 bits per color) and restored images using our method. Best viewed zoomed in.

Figure 3: Rate-distortion curves for standard JPEG (blue) and our method (green).

To address this issue, we introduce a method that performs JPEG artifact correction using Denoising Diffusion Restoration Models (DDRM) [16]. Our core idea is to generalize the pseudo-inverse matrix that exists in the DDRM update rule for the noiseless observation case. This generalized notion of a “pseudo-inverse” includes JPEG as a special case, where the “pseudo-inverse” for JPEG encoding is simply JPEG decoding. The resulting algorithm resembles the original update for DDRM, replacing the linear operator and its pseudo-inverse with JPEG encoding and decoding, respectively.

We apply our algorithm for JPEG restoration with various quality factors (QF), where the quantization matrices are embedded in the JPEG files and naturally known at inference time. In common image quality metrics such as PSNR, SSIM [38], and LPIPS [41], our method compares favorably against a recent state-of-the-art GAN-based baseline [40] trained specifically for JPEG restoration. Our method achieves even more improvement on low QF that the baseline is not trained on, demonstrating the generalization advantages of methods that leverage unconditional diffusion models.

2 Background

Diffusion Models. Diffusion models [31, 13] are generative models with a Markov chain structure $x_T \rightarrow x_{T-1} \rightarrow \dots \rightarrow x_1 \rightarrow x_0$; where $x_t \in \mathbb{R}^n$, which defines the following joint distribution:

$$p(x_{0:T}) = p^{(T)}(x_T) \prod_{t=0}^{T-1} p^{(t)}(x_t | x_{t+1});$$

After x_T, \dots, x_0 are drawn, only x_0 is kept as the final output of the generative model. When training a diffusion model, a fixed, factorized variational inference distribution is assumed:

$$q(x_{1:T} | x_0) = q^{(T)}(x_T | x_0) \prod_{t=0}^{T-1} q^{(t)}(x_t | x_{t+1}; x_0);$$

which leads to an evidence lower bound (ELBO) on the maximum likelihood objective. One particular parametrization [32] takes the following form:

$$q(x_t | x_0) = N(p_{-t} x_0; (1 - \beta_t) I); \quad \beta_t \in [0, 1] \quad (1)$$

where $p^{(t)}$ can be trained via a denoising autoencoder [36] objective; in the ideal case, the denoiser, denoted $p^{(t)}(x_t)$ should map to the MMSE (Minimum Mean Squared Error) estimator $E_{q(x_0|x_t)}[x_0]$, and constitutes a prediction over the “clean”. Diffusion models have achieved unprecedented success in image generation [35, 8, 17], and they have also been deployed for a variety of tasks [19, 6, 25, 3, 12, 1, 2, 29, 30].

Linear Inverse Problems. A general linear inverse problem is posed as

$$y = Hx + z; \quad (2)$$

where our aim is to recover the signal $x \in \mathbb{R}^n$ from measurements $y \in \mathbb{R}^m$. $H \in \mathbb{R}^{m \times n}$ is a known degradation matrix, and $z \sim N(0; \frac{\sigma^2}{\gamma} I)$ is an additive white Gaussian noise with known variance.

Various works have applied diffusion models for inverse problem solving, mostly for the noiseless case. While it is possible to train a conditional diffusion model based on pairs only [27, 28, 39], such models may not generalize to other inverse problems. Therefore, it is often desirable to formulate inverse problem solvers from unconditional diffusion models [5, 34, 18, 16], where the knowledge about the inverse problem does not need to be known during training; compared with problem-specific conditional diffusion models, problem-agnostic techniques save significant computational resources.

Table 1: JPEG artifact correction results on ImageNet-1K for varying quality factors (QF).

Method	QF = 5			QF = 10			QF = 30		
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
JPEG	23:46	0:62	0:48	26:12	0:73	0:35	29:53	0:84	0:19
QGAC [10]	23:85	0:64	0:43	28:01	0:80	0:26	31:28	0:89	0:15
Ours (S)	25:19	0:70	0:34	27:68	0:78	0:26	30:98	0:87	0:15
Ours (A)	25:90	0:72	0:34	28:37	0:80	0:25	31:58	0:88	0:15

Denoising Diffusion Restoration Models (DDRM). In particular, DDRM [16] is a general solver for linear inverse problems in both noisy and noiseless cases. For any linear inverse problem, the DDRM model is defined as

$$p(x_{0:T}|y) = p^{(T)}(x_T|y) \prod_{t=0}^{T-1} p^{(t)}(x_t|x_{t+1}; y);$$

where x_0 is the final diffusion output. The high-level idea behind DDRM is to leverage the singular value decomposition of H and transform both x and the possibly noisy y to a shared spectral space. In this space, DDRM performs denoising on dimensions when information from y is available (i.e., when singular values are non-zero) and performs imputation on ones where such information is not available (i.e., when singular values are zero), taking account of the measurement noise explicitly.

JPEG. JPEG [8] is a commonly used lossy compression method for images. At a high-level, JPEG first transforms an uncompressed image from the RGB color space to the YCbCr space, optionally applies chroma subsampling, splits the image into 8 pixel blocks, performs a discrete cosine transform (DCT), and then performs quantization of the resulting values using a fixed quantization matrix. These values can then be compressed in a lossless fashion via Huffman trees. The entire process can be reverted to define the JPEG decoding method, with loss of information happening in the chroma subsampling and the quantization steps. Since its introduction in 1991, JPEG has become the most widely used image compression format in the world, with several billion JPEG images produced every day; thus, restoring high-quality images from JPEG-compressed ones has wide applications. Several previous methods were developed for this purpose [10, 9, 23, 10, 15].

3 JPEG Artifact Correction with DDRM

For the case of no noise in the observation, the general DDRM process to sample from $p^{(t)}(x_t|x_{t+1}; y)$ for linear inverse problems simplifies to be

$$x_t^0 = f^{(t+1)}(x_{t+1}) - H^y H f^{(t+1)}(x_{t+1}) + H^y y;$$

$$x_t = p_{t-b} x_t^0 + (1 - p_{t-b}) f^{(t+1)}(x_{t+1}) + p_{t+1} \frac{1}{1 - p_{t+1}} (1 - p_{t+1})^{(t+1)}(x_{t+1})$$

where H^y is the Moore-Penrose pseudo-inverse of H , $f^{(t+1)}(x_{t+1})$ is the denoising model output at the previous step $t+1$, $(1 - p_{t-b}) f^{(t+1)}(x_{t+1})$ is the predicted noise value, and p_{t-b} are user-defined hyperparameters, and $N(0; I)$ is a standard Gaussian distributed vector. At a high level, we inject information about y via x_t^0 where we replace the values in the spectral domain with what we know from y . The sampling process then aggregates x_t^0 (the current input), and $f^{(t+1)}(x_{t+1})$ (the denoiser output) to produce the value for the next iteration.

While it seems that the above approach only works for linear inverse problems, its insights can actually be used for other non-linear inverse problems, such as JPEG artifact correction. We note that for a linear pseudo-inverse H^y has two important properties:

1. $HH^yH = H$, i.e., taking the pseudo-inverse does not change the measurement.
2. H^yHx is "close" to x , in the sense that $\|H^yHx - x\|_2^2$ provides a least squares solution for the problem $\min_z \|kz - x\|_2^2$ for all x when observing only $y = Hx$ (but not x).

The above properties may exist for operators that are not linear. For example, if we treat the JPEG encoding operator, then the JPEG decoding operator also satisfies these properties:

1. The JPEG encoding introduces loss of information during the quantization and chroma subsampling stages. The remaining information is kept during JPEG decoding, and thus encoding it again will lead to the same result.
2. The JPEG decoding method generally preserves visual similarity, so applying decoding after encoding should generate an image that is “close” to the original one.

With this insight, we can simply perform JPEG restoration with DDRM with the update rule

$$x_t^0 = f^{(t+1)}(x_{t+1}) \text{ Decode Encode } f^{(t+1)}(x_{t+1}) + \text{Decode}(y);$$

$$x_t = \frac{p}{t} x_t^0 + (1 - \frac{p}{t}) f^{(t+1)}(x_{t+1}) + \frac{p}{t} \frac{1}{t} x_t + (1 - \frac{p}{t}) f^{(t+1)}(x_{t+1});$$

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

4 Experimental Results

We evaluate our method on the ImageNet dataset, as it is diverse and represents real-world use cases. Specifically, we evaluate on a 1000-image subset of the ImageNet validation set named ImageNet-1K [26]. We utilize the diffusion model from [8], trained on 256 × 256-pixel ImageNet training images, for a diffusion schedule of 1000 timesteps. In all of our experiments, we choose the hyperparameters $\beta_1 = 1$, $\beta_b = 0.4$, and 20 uniformly-spaced diffusion steps. Additionally, since JPEG images generally preserve overall image contents, we find that we can perturb a JPEG-compressed image with noise and use it as an initialization for our sampling process at an intermediate step $t = 300$, similar to [24]. This allows the sampling to provide more faithful reconstructions, avoiding the unnecessary randomness induced by starting at the initial timestep $t = 1000$. However, as we use a probabilistic sampling scheme, randomness can still be expected in the results. In order to stabilize performance, we draw 30 independent samples for each input and save the resulting average image. We denote the first sample “ours (S)” and the averaged image “ours (A)”.

For our JPEG artifact correction experiments, we use the most common variant of JPEG which includes chroma subsampling and quantization matrices defined by a quality factor (QF) ranging from 1 to 100, 1 being the most compressed, and 100 being the most faithful to the original image. Our method produces high quality reconstructions (see Figures 1, 4). Moreover, when evaluated numerically on common metrics such as PSNR, SSIM, and LPIPS [41], our method provides a significant improvement over simple JPEG decoding, and its performance is favorable or comparable to a recent state-of-the-art JPEG artifact correction named QGAC [38]. QGAC was specifically trained for JPEG restoration with QF [10; 100], and as can be seen in Table 1, it generalizes poorly for lower QF. In contrast, our method generalizes well for all QF without JPEG-specific training. We demonstrate its success by showing its compression rate-distortion curve in Figure 3.

Furthermore, our method is not limited to JPEG artifact correction, but can also be applied to similar non-linear inverse problems not covered by DDRM [6]. For instance, we consider the problem of image dequantization, where we attempt to recover high-quality reconstructions from images that were quantized below the standard 24 bits per color. As evident in Figure 2, our method generalizes well for image dequantization, owing to its problem-agnostic nature.

5 Conclusion

We propose a novel method for correcting JPEG compression artifacts using diffusion models. Our method extends DDRM [6] beyond the linear case by generalizing the pseudo-inverse concept. We perform evaluations on ImageNet-1K [26] where our method performs on par with a state-of-the-art baseline in most cases, and exhibits generalization abilities for lower quality factors which specifically trained baselines do not possess. Our method can further generalize beyond JPEG restoration, as we successfully demonstrate on the image dequantization problem. It does so seamlessly, without retraining, and without problem-specific hyperparameter tuning.

References

- [1] Tomer Amit, Eliya Nachmani, Tal Shaharabany, and Lior Wolf. Segdiff: Image segmentation with diffusion probabilistic models. *arXiv preprint arXiv:2112.00390*, 2021.
- [2] Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18208–18218, 2022.
- [3] Tsachi Blau, Roy Ganz, Bahjat Kawar, Alex Bronstein, and Michael Elad. Threat model-agnostic adversarial defense using diffusion models. *arXiv preprint arXiv:2207.08089*, 2022.
- [4] Alfred M Bruckstein, Michael Elad, and Ron Kimmel. Down-scaling for better transform compression. *IEEE Transactions on Image Processing*, 12(9):1132–1144, 2003.
- [5] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. Ilvr: Conditioning method for denoising diffusion probabilistic models. *arXiv preprint arXiv:2108.02938*, 2021.
- [6] Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, 34:17695–17709, 2021.
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009.
- [8] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- [9] Chao Dong, Yubin Deng, Chen Change Loy, and Xiaoou Tang. Compression artifacts reduction by a deep convolutional network. *Proceedings of the IEEE international conference on computer vision*, pages 576–584, 2015.
- [10] Max Ehrlich, Larry Davis, Ser-Nam Lim, and Abhinav Shrivastava. Quantization guided JPEG artifact correction. In *European Conference on Computer Vision*, pages 293–309. Springer, 2020.
- [11] Independent JPEG Group. Libjpeg, 1991.
- [12] Xizewen Han, Huangjie Zheng, and Mingyuan Zhou. CARD: Classification and regression diffusion models. *arXiv preprint arXiv:2206.07275*, 2022.
- [13] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, volume 33, pages 6840–6851, 2020.
- [14] Ajil Jalal, Marius Arvinte, Giannis Daras, Eric Price, Alex Dimakis, and Jonathan Tamir. Robust compressed sensing mri with deep generative priors. *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- [15] Jiayi Jiang, Kai Zhang, and Radu Timofte. Towards visible blind jpeg artifacts removal. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4997–5006, 2021.
- [16] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. In *Advances in Neural Information Processing Systems*, 2022.
- [17] Bahjat Kawar, Roy Ganz, and Michael Elad. Enhancing diffusion-based image synthesis with robust classifier guidance. *arXiv preprint arXiv:2208.08664*, 2022.
- [18] Bahjat Kawar, Gregory Vaksman, and Michael Elad. SNIPS: Solving noisy inverse problems stochastically. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021.
- [19] Bahjat Kawar, Gregory Vaksman, and Michael Elad. Stochastic image denoising by sampling from the posterior distribution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 1866–1875, 2021.

- [20] Orest Kupyn, Tetiana Martyniuk, Junru Wu, and Zhangyang Wang. Deblurgan-v2: Deblurring (orders-of-magnitude) faster and better. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8878–8887, 2019.
- [21] Anders Boesen Lindbo Larsen, Søren Kaae Sønderby, Hugo Larochelle, and Ole Winther. Autoencoding beyond pixels using a learned similarity metric. *International conference on machine learning*, pages 1558–1566. PMLR, 2016.
- [22] Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-Realistic single image Super-Resolution using a generative adversarial network. preprint arXiv:1609.04802, September 2016.
- [23] Pengju Liu, Hongzhi Zhang, Kai Zhang, Liang Lin, and Wangmeng Zuo. Multi-level wavelet-cnn for image restoration. *Proceedings of the IEEE conference on computer vision and pattern recognition workshop*, pages 773–782, 2018.
- [24] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. SDEdit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- [25] Weili Nie, Brandon Guo, Yujia Huang, Chaowei Xiao, Arash Vahdat, and Anima Anandkumar. Diffusion models for adversarial purification. *International Conference on Machine Learning (ICML)*, 2022.
- [26] Xingang Pan, Xiaohang Zhan, Bo Dai, Dahua Lin, Chen Change Loy, and Ping Luo. Exploiting deep generative prior for versatile image restoration and manipulation. *European Conference on Computer Vision (ECCV)*, 2020.
- [27] Chitwan Saharia, William Chan, Huiwen Chang, Chris A Lee, Jonathan Ho, Tim Salimans, David J Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. preprint arXiv:2111.05826, 2021.
- [28] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. preprint arXiv:2104.07636, 2021.
- [29] Yuyang Shi, Valentin De Bortoli, George Deligiannidis, and Arnaud Doucet. Conditional simulation using diffusion Schrödinger bridges. preprint arXiv:2202.13460, 2022.
- [30] Abhishek Sinha, Jiaming Song, Chenlin Meng, and Stefano Ermon. D2c: Diffusion-decoding models for few-shot conditional generation. *Advances in Neural Information Processing Systems*, 34:12533–12548, 2021.
- [31] Jascha Sohl-Dickstein, Eric A Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. preprint arXiv:1503.03585, March 2015.
- [32] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, April 2021.
- [33] Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. preprint arXiv:2111.08005, 2021.
- [34] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. preprint arXiv:2011.13456, 2020.
- [35] Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. In *Neural Information Processing Systems (NeurIPS)*, 2021.
- [36] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.

- [37] Gregory K Wallace. The JPEG still picture compression standard. *Communications of the ACM* 34(4):30–44, 1991.
- [38] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13(4):600–612, 2004.
- [39] Jay Whang, Mauricio Delbracio, Hossein Talebi, Chitwan Saharia, Alexandros G Dimakis, and Peyman Milanfar. Deblurring via stochastic re-norming. *arXiv preprint arXiv:2112.02475* 2021.
- [40] Zixiang Xiong, Michael T Orchard, and Ya-Qin Zhang. A deblocking algorithm for jpeg compressed images using overcomplete wavelet representation. *IEEE Transactions on Circuits and Systems for Video Technology* 7(2):433–437, 1997.
- [41] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. *Proceedings of the IEEE conference on computer vision and pattern recognition* pages 586–595, 2018.

A Additional Visual Results




Figure 4: Triplets of original (ground-truth), JPEG compressed, and restored images. Across different quality factors (QF), our method successfully corrects artifacts of JPEG compression. Images are accompanied by a zoomed-in area in the bottom right corner to highlight specific artifact removals.