

OPENDENOISING: AN EXTENSIBLE BENCHMARK FOR BUILDING COMPARATIVE STUDIES OF IMAGE DENOISERS

Florian Lemarchand, Eduardo Fernandes Montesuma, Maxime Pelcat, Erwan
Nogues

► **To cite this version:**

Florian Lemarchand, Eduardo Fernandes Montesuma, Maxime Pelcat, Erwan Nogues. OPENDENOISING: AN EXTENSIBLE BENCHMARK FOR BUILDING COMPARATIVE STUDIES OF IMAGE DENOISERS. ICASSP 2020-IEEE International Conference on Acoustics, Speech, and Signal Processing, May 2020, Barcelone, Spain. hal-02477648

HAL Id: hal-02477648

<https://hal.archives-ouvertes.fr/hal-02477648>

Submitted on 13 Feb 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

OPENDENOISING: AN EXTENSIBLE BENCHMARK FOR BUILDING COMPARATIVE STUDIES OF IMAGE DENOISERS

Florian Lemarchand ^{*}, Eduardo Fernandes Montesuma ^{*}, Maxime Pelcat ^{*}, Erwan Nogues ^{*§}

^{*} Univ. Rennes, INSA Rennes, IETR - UMR CNRS 6164

[§] DGA-MI, Bruz

ABSTRACT

Image denoising has recently taken a leap forward due to machine learning. However, image denoisers, both expert-based and learning-based, are mostly tested on well-behaved generated noises (usually Gaussian) rather than on real-life noises, making performance comparisons difficult in real-world conditions. This is especially true for learning-based denoisers which performance depends on training data. Hence, choosing which method to use for a specific denoising problem is difficult.

This paper proposes a comparative study of existing denoisers, as well as an extensible open tool that makes it possible to reproduce and extend the study. MWCNN is shown to outperform other methods when trained for a real-world image interception noise, and additionally is the second least compute hungry of the tested methods. To evaluate the robustness of conclusions, three test sets are compared. A Kendall’s Tau correlation of only 60% is obtained on methods ranking between noise types, demonstrating the need for a benchmarking tool.

Index Terms— Image Denoiser Benchmarking, Complex Image Noises

1. INTRODUCTION

Image denoising, as a sub-domain of image restoration, is an extensively studied problem. The objective of a denoiser is to restore a cleaned image signal $\hat{\mathbf{x}}$ from an observation \mathbf{y} considered to be a noisy or corrupted version of an original image \mathbf{x} . \mathbf{y} is generated by a noise function \mathbf{H} such that $\mathbf{y} = \mathbf{H}(\mathbf{x})$. If the noise is additive, \mathbf{H} is a sum and $\mathbf{y} = \mathbf{x} + \eta$, where η is a realisation of \mathbf{H} . There exist plenty of noise models to represent \mathbf{H} among which examples of frequently used models are additive Gaussian, Poisson, salt and pepper or speckle noises. State of the art denoisers are constantly progressing in terms of noise elimination level [1, 2, 3]. However, most techniques are tailored for and evaluated on a given noise distribution, exploiting its probabilistic properties to distinguish it from the signal of interest. On the specific case of Gaussian additive noise, current denoisers are approaching theoretical bounds [4].

Besides the largely addressed *well-behaved* noise models, image restoration is also concerned by more complex noise distributions. While these distributions are application specific, they are real-world cases directly issued from identified technical needs such as image interception in difficult conditions.

In a context where new methods are constantly appearing, it is challenging to fairly compare emerging methods to previous ones. Moreover, when a real-world noise needs to be eliminated, it is difficult to determine which of method is the best for the given noise characteristics. Even if most state of the art methods are evaluated on the de-facto standard databases (e.g. BSD [5], 12 well-known

images as Lenna, Cameraman, etc.), methods addressing specific noises and image types have to be evaluated on tailored databases.

In this context, the contributions of this paper are:

- An extensible and open-source benchmark for comparing image restoration methods.
- A comparative study of current denoisers on mixture and interception noise elimination, as a use case for the benchmark.

The paper is organized as follows. Section 2 presents state of the art methods for image denoising as well as existing solutions to benchmark them. Section 3 describes the proposed benchmark. The comparative study, covering six restoration methods, is proposed in Section 4. Section 5 concludes the paper.

2. RELATED WORK

2.1. Related Work on Benchmarks of Image Denoisers

An active research domain in complex noise restoration is photograph restoration. This domain aims at removing a noise introduced by sensor hardware defects. Supervised datasets can be built by calibrating a sensor and hence obtaining pairs of clean and noisy samples. Darmstadt [6] and PolyU [7] datasets are such datasets. Authors propose to use their datasets as a means for benchmarking denoising algorithms. This work is complementary to our proposed benchmark that can adapt to different datasets.

Open-source projects have been created to benchmark denoising methods. The University of Toronto proposes a benchmark ¹ to provide reproducibility for a method proposed in 2010 [8]. This benchmark is tailored to the solution and not built to be extended. Another unpublished benchmark exists that implements denoising as well as other restoration algorithms such as super-resolution ². The benchmark is limited to learning-based, Python-implemented and pre-trained methods. The latter limitation drastically reduces the use of such benchmark for complex noises. Indeed, most state of the art methods, when delivered trained, are trained on well-behaved noise. This paper proposes a benchmark extensible in several aspects. Indeed, the user can introduce his datasets, denoising methods, and metrics.

2.2. Related Work on Image Denoisers

Image denoising techniques are as old as image sensors whose defects they counteract. Current denoising solutions are either expert-based denoisers, human crafted based on an expertise of artifacts or

¹www.cs.utoronto.ca/~strider/Denoise/Benchmark/

²<https://github.com/titsitits/open-image-restoration>




Fig. 1. OpenDenoising Block diagram. Users can tune datasets, metrics, denoising models and evaluation functions. OpenDenoising produces denoised samples as well as performance metrics.

of statistical noise properties, or learning-based denoisers leveraging on latent image priors extracted from data.

Block-Matching 3D (BM3D) [1] is a state-of-the-art expert-based method for Additive White Gaussian Noise (AWGN) removal. When introduced in 2007, BM3D clearly outperformed the previous state of the art TNRD method by a good margin of 2dB. BM3D performs block matching to find patches with similar content in the image and uses collaborative filtering, thresholding and Wiener filtering into the transform domain to restore the image.

Authors of [9] were the firsts to propose an encoding/decoding model with denoising criterion. Their proposal named stacked auto-encoder learns to map the noisy image to a latent space (encoding) and projects back the latent representation to the input space (decoding) to produce the denoised image. More recent auto-encoders have been proposed such as RED [10] that adds convolutional layers and uses skip connections to better keep image priors throughout the encoding/decoding pipeline.

Following these premises of denoising auto-encoders, several Convolutional Neural Network (CNN) methods have emerged such as Denoising Convolutional Neural Network (DnCNN) [2]. DnCNN exploits residual learning, i.e. it learns to isolate the noise \mathbf{H} from the corrupted sample to later remove this noise instead of directly recovering the latent clean signal. Multi-level Wavelet Convolutional Neural Network (MWCNN) [3] is also CNN-based. Its novelty lies in the symmetrical use of wavelet and inverse wavelet transforms into the contracting and expanding parts of a U-Net [11] architecture. The use of wavelet enables *safe* subsampling with no information loss providing a better recovering of textures and sharp structures.

The most recent learning-based methods are less supervised, i.e. they require less noisy/clean image pairs to train. In Noise2Void [12] and Noise2Noise[13], authors propose a tactic to train a denoising model using a single noise realisation. Authors introduce the idea of *blind-spot masking* during training. They claim that the essential advantage of that strategy is to avoid to learn the identity due to the masking of the central value of the receptive field.

These methods are evaluated on well-behaved noises (typically AWGN) for the tests to be easily reproducible. Only Noise2Void is evaluated on medical images subject to complex noise. In the following, our open benchmark is proposed to assess denoisers' quality.

3. PROPOSED BENCHMARK

Considering the above discussed issues with existing benchmarks, we propose the OpenDenoising benchmark illustrated in Fig. 1. It is an open-source tool with tutorials and documentation³ released under a CeCILL-C license. OpenDenoising is implemented in Python and has been designed for extensions. Adding a new denoiser to

³<https://github.com/opendenoising/benchmark>




Fig. 2. Inference time (log scale) for different denoisers. Image resolution is 256×256 . Noise2Void is the fastest method by almost 10 folds. MWCNN, RED10, DnCNN and RED30 are close to each other. BM3D is the slowest with an inference time over one second. Setup: Intel Xeon W-2125 CPU and Nvidia GTX1080 Ti GPU.

the benchmark is a matter of minutes following a tutorial and opens for comparison with the built-in methods evoked in Section 4. For learning-based methods, the application is compatible and tested with most major frameworks (Tensorflow, Keras, Pytorch, Matlab). For learning-based training and evaluation, it is possible to use one or several datasets either supervised or not. Any scalar metric being coded in Python can be used in the benchmark. Several preprocessing functions, e.g. for data augmentation, are provided, and custom functions can be introduced.

The user chooses whether a training is required for a method and in that case selects training parameters. Once the training is launched, monitors can be output by the benchmark to observe the learning phase. When trained models are available, evaluation is launched with custom or built-in metrics. The results are outlined using custom or built-in plots and/or stored as images or csv summaries.

As an example of OpenDenoising versatility, it would be possible to extend the benchmark to classification methods only implementing custom evaluations metrics. Other potential usages of OpenDenoising include: study the extensibility of methods to new applications (see Section 4), study the strategies for re-training off-the-shelf methods (from scratch or with fine-tuning), and tune hyperparameters. The experimental results presented in the next section exploit OpenDenoising to build a comparative study of state of the art denoisers on various types of noise.

4. A COMPARATIVE STUDY OF DENOISERS

In this section, we apply top-ranking denoisers to images with various noises. For comparison fairness of training-based methods, no data augmentation is made and the same training datasets are used. This setup excepted, methods are trained (when applicable) using original papers parameters and training strategies. Four noise types with increasing complexity are exploited to observe the behavior of the studied denoisers. Peak Signal to Noise Ratio (PSNR) and Structural Similarity (SSIM) are used to respectively evaluate the objective and subjective quality of the denoised image.

4.1. Gaussian Noise

First, Gaussian noise is denoised to test the methods in their original conditions. Denoisers are evaluated on a common noisy dataset corrupted with AWGN. The underlying data is made of 10k natural images extracted from the ImageNet [14] evaluation set. Average PSNR and SSIM are shown in Table 1 and example images displayed

	Dataset	No Denoising	BM3D	RED10	RED30	DnCNN-B	MWCNN	Noise2Void
PSNR	Gaussian	14.96	23.90	25.52	25.82	25.67	25.49	23.41
	Mixture	10.58	18.29	24.25	24.58	24.50	24.30	19.93
	Interception-Like	17.16	22.04	51.56	52.08	51.66	51.16	21.70
	Interception	9.46	9.61	22.59	23.46	23.04	23.66	9.46
SSIM	Gaussian	0.24	0.67	0.71	0.73	0.72	0.72	0.62
	Mixture	0.12	0.31	0.67	0.68	0.68	0.68	0.51
	Interception-Like	0.11	0.98	0.99	0.99	0.99	0.99	0.98
	Interception	0.32	0.73	0.94	0.96	0.95	0.96	0.47

Table 1. Average evaluation of PSNR and SSIM metrics on test sets for, from top to bottom row: AWGN with noise level $\sigma = 50$; Mixture noise made of AWGN with noise level $\sigma = 50$ and salt and pepper 20% ; Interception-Like noise being interception reference samples noised with AWGN with noise level $\sigma = 50$; Interception noise. The test set for each noise is made of 200 samples excluded from training set.




Fig. 3. PSNR of denoised images on (a) Gaussian noise, (b) Mixture and (c) Interception noise. Outliers are not displayed.

on Fig. 4. Fig. 3 shows in boxplots the 10th, 25th, 75th and 90th percentiles of PSNR results as well as the median. To focus on difficult noises, the maximum noise level commonly found in papers is picked, namely $\sigma = 50$. Experimental results, shown on the first line of Table 1, are coherent with the published ones, though slightly under because no data augmentation is applied. On Gaussian noise, RED30 outperforms other methods (by a limited 0.15dB Δ PSNR and 1% Δ SSIM) but it is also the most costly in terms of number of parameters.

4.2. Mixture Noise

Complicating the denoising task, a mixture noise is then studied. This mixture noise is constructed through the successive corruption of the samples by the previously used AWGN ($\sigma = 50$) and an additional salt and pepper corruption (20% of corrupted pixels, half 0, half maximum). The clean data is conserved and this noise mixture roughly models the behaviour of an image sensor introducing Gaussian noise because of its hardware non-uniformity and salt and

pepper due to pixel defects.

Fig. 3 shows that learning-based methods perform consistently better than BM3D. BM3D is here used out of its original objective (i.e. Gaussian denoising) and thus performs poorly. Another information brought by mixture noise is that Noise2Void clearly underperforms compared to other learning-based methods. This is not surprising considering the addition of salt and pepper noise that damages the spatial coherence used as a hypothesis in the Noise2Void strategy. RED10, RED30, DnCNN and MWCNN have close performances with a narrow victory for RED30 (0.08dB Δ PSNR).

4.3. Interception Noise

A real-world complex noise is now studied, generated by intercepting images from Electro Magnetic (EM) emanations. Electronic devices produce EM emanations that not only interfere with radio devices but also compromise the data they handle. A third party performing a side-channel analysis can recover internal information from both analog [15], and digital circuits [16]. Following an eavesdropping procedure [17], it is possible to build a supervised dataset made of pairs of reference images, originally displayed on a screen, and their intercepted noisy versions. Interception strongly damages the images and denoising is necessary to interpret their content. For reproducibility, we released the dataset used for this study⁴. It contains more than 120k samples.

To study the noise complexity, the intercepted clean samples are also artificially corrupted using AWGN with $\sigma = 50$. The resulting samples are called interception-like. As shown in Table 1, most methods perform well on that denoising task. The clean content of the intercepted samples contains black characters printed on a white background. The latent clean distribution of samples is thus not an issue to denoise with learning-based methods, Noise2Void excluded. Only BM3D and Noise2Void have problems with background restoration (see Fig. 4). This phenomena is due to the correlation of the samples content and the noising process. The background of the clean samples is fully white. When applying AWGN, half of the noise coefficients are negative and samples are clipped to an integer format. Thus, the assumption of a Gaussian distribution does not hold, leading to poor restoration results with non-supervised methods, unable to adapt.

Table 1, Fig. 3 and Fig. 4 show the results of denoising methods applied to interception noise. Metrics drop for all methods on this complex noise. Noise2Void does not manage to denoise at all. As explained in the original paper, Noise2Void has difficulties with noise correlated between several pixels which is here the case. BM3D is

⁴https://github.com/opendenoising/interception_dataset




Fig. 4. From top to bottom, one sample per dataset is shown noisy (left), denoised with different denoisers (middle) and clean (right). PSNR/SSIM are displayed for each sample. Images with the best compromise between PSNR and SSIM metrics are yellow-boxed.

built for AWGN and is not trainable, hence the poor results in that case. Others learned methods like RED10 and DnCNN produce interesting denoising but perceptual results of Fig. 4 show hardly interpretable samples, not revealed by SSIM. RED30 and MWCNN are the best-performing methods for interception noise removal but still with some extant artifacts.

4.4. Discussion

Different conclusions can be drawn from the above experiments using OpenDenoising and 4 different datasets.

First, Fig. 3 shows that the performance ranking between methods strongly depends on noise type, training dataset and evaluation dataset. When calculating *Kendall's Tau correlation coefficient* [18], a value of 60% is obtained between Gaussian noise ranking and Interception noise ranking, ranking based on the mean PSNR. This correlation coefficient, while being high - Kendall's Tau is in $[-1; 1]$; -1 and 1 respectively meaning fully discordant and fully concordant rankings -, shows the need for a benchmark such as the proposed. It automates the comparison process and the selection of a given method for a denoising problem. As an example, it would be a wrong choice to pick RED30 instead of MWCNN for interception restoration based on the original paper evaluations (e.g. Gaussian evaluation). MWCNN is indeed both more efficient and less computationally intensive than RED30, as shown in Fig. 2.

Results of Table 1, Fig. 3 and Fig. 4 show a growing gap between expert-based and learning-based method as the complexity of the denoising increase. This can be explained by the flexibility of learning-based models and the advanced information brought by a supervised training. This is evidenced by the low performance of the non-supervised Noise2Void on Mixture and Interception noises.

As stated in several studies, PSNR and SSIM do not suit well the assessment of interception restoration. Fig. 4 shows that

while evaluation metrics on interception noise are reasonably good (PSNR/SSIM values around 20dB/0.9), the perceptual quality is poor. The explanation lies in the latent content of the samples, made of black characters on a white background. A good background restoration is sufficient to raise good evaluation metrics. This issue is evoked in [19] where authors propose a different evaluation metric to overcome the problem.

5. CONCLUSIONS

In this paper, the OpenDenoising tool has been proposed. OpenDenoising benchmarks image denoisers and aims at comparing methods on a common ground in terms of datasets, training parameters and evaluation metrics. Supporting several languages and learning frameworks, OpenDenoising is also extensible and open-source. The second contribution of the paper is a comparative study of image restoration in the case of a complex noise source.

Three major conclusions arise from the comparative study. First, the difference in terms of performance between expert-based and learning-based methods rises as the complexity of the noise grows. Second, the ranking of methods is strongly impacted by the nature of the noises. Finally, MWCNN proves to be the best method for the considered real-world interception restoration task. It slightly outperforms DnCNN and RED30 while being substantially faster.

These results show that restoring an image from a complex noise is not universally solved by a single method and that choosing a denoiser requires automated testing. Our future work includes the design of new denoising methods for intercepted images.

6. ACKNOWLEDGEMENTS

This work is supported by the "Pôle d'Excellence Cyber", initiative of the French Ministry of the Armed Forces and Bretagne council.

7. REFERENCES

- [1] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image Denoising by Sparse 3-D Transform-Domain Collaborative Filtering," *IEEE Transactions on Image Processing*, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.
- [2] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a Gaussian Denoiser: Residual Learning of Deep CNN for Image Denoising," *IEEE Transactions on Image Processing*, vol. 26, no. 7, pp. 3142–3155, July 2017.
- [3] P. Liu, H. Zhang, K. Zhang, L. Lin, and W. Zuo, "Multi-level Wavelet-CNN for Image Restoration," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, Salt Lake City, UT, June 2018, pp. 886–88609, IEEE.
- [4] P. Chatterjee and P. Milanfar, "Is Denoising Dead?," *IEEE Transactions on Image Processing*, vol. 19, no. 4, pp. 895–911, Apr. 2010.
- [5] D. Martin, C. Fowlkes, D. Tal, and J. Malik, "A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics," in *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, Vancouver, BC, Canada, 2001, vol. 2, pp. 416–423, IEEE Comput. Soc.
- [6] T. Plotz and S. Roth, "Benchmarking Denoising Algorithms with Real Photographs," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI, July 2017, pp. 2750–2759, IEEE.
- [7] J. Xu, H. Li, Z. Liang, D. Zhang, and L. Zhang, "Real-world Noisy Image Denoising: A New Benchmark," *arXiv:1804.02603 [cs]*, Apr. 2018, arXiv: 1804.02603.
- [8] F. Estrada, D. Fleet, and A. Jepson, "Stochastic Image Denoising," in *Proceedings of the British Machine Vision Conference 2009*, London, 2009, pp. 117.1–117.11, British Machine Vision Association.
- [9] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.A. Manzagol, "Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion," *Journal of Machine Learning Research*, vol. 11, pp. 3371–3408, 2010.
- [10] X. Mao, C. Shen, and Y-B. Yang, "Image Restoration Using Very Deep Convolutional Encoder-Decoder Networks with Symmetric Skip Connections," *Advances in Neural Information Processing Systems 29 (NIPS 2016)*, p. 9, 2016.
- [11] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*. 2015, Lecture Notes in Computer Science, pp. 234–241, Springer.
- [12] A. Krull, T-O. Buchholz, and F. Jug, "Noise2void - Learning Denoising From Single Noisy Images," *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, p. 9, 2019.
- [13] J. Lehtinen, J. Munkberg, J. Hasselgren, S. Laine, T. Karras, M. Aittala, and T. Aila, "Noise2noise: Learning Image Restoration without Clean Data," *CoRR*, 2018.
- [14] J. Deng, W. Dong, R. Socher, L. Li, Kai L., and Li F-F., "ImageNet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, Miami, FL, June 2009, pp. 248–255, IEEE.
- [15] W. Van Eck, "Electromagnetic radiation from video display units: An eavesdropping risk?," *Computers & Security*, vol. 4, no. 4, pp. 269–286, 1985.
- [16] M. G. Kuhn, "Compromising Emanations of LCD TV Sets," *IEEE Transactions on Electromagnetic Compatibility*, vol. 55, no. 3, pp. 564–570, 2013.
- [17] F. Lemarchand, C. Marlin, F. Montreuil, E. Nogues, and M. Pelcat, "Electro-Magnetic Side-Channel Attack Through Learned Denoising and Classification," *arXiv:1910.07201 [cs, eess]*, Oct. 2019, arXiv: 1910.07201.
- [18] Haris Javaid, Aleksander Ignjatovic, and Sri Parameswaran, "Fidelity metrics for estimation models," in *2010 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, San Jose, CA, USA, Nov. 2010, pp. 1–8, IEEE.
- [19] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual Losses for Real-Time Style Transfer and Super-Resolution," in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds., vol. 9906, pp. 694–711. Springer International Publishing, Cham, 2016.