Regularizing the Deep Image Prior with a Learned Denoiser for Linear Inverse Problems

Rita Fermanian, Mikael Le Pendu, Christine Guillemot

▶ **To cite this version:**

Rita Fermanian, Mikael Le Pendu, Christine Guillemot. Regularizing the Deep Image Prior with a Learned Denoiser for Linear Inverse Problems. MMSP 2021 - IEEE 23rd International Workshop on Multimedia Siganl Processing, Oct 2021, Tampere, Finland. pp.1-6. hal-03310533

HAL Id: hal-03310533
https://hal.science/hal-03310533v1

Submitted on 30 Jul 2021

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.

Regularizing the Deep Image Prior with a Learned Denoiser for Linear Inverse Problems

Rita Fermanian, Mikael Le Pendu, Christine Guillemot
Inria Rennes – Bretagne-Atlantique, 263 Avenue Général Leclerc, 35042 Rennes Cedex, France
(e-mail: firstname.lastname@inria.fr).

***Abstract*—We propose an optimization method coupling a learned denoiser with the untrained generative model, called deep image prior (DIP) in the framework of the Alternating Direction Method of Multipliers (ADMM) method. We also study different regularizers of DIP optimization, for inverse problems in imaging, focusing in particular on denoising and super-resolution. The goal is to make the best of the untrained DIP and of a generic regularizer learned in a supervised manner from a large collection of images. When placed in the ADMM framework, the denoiser is used as a proximal operator and can be learned independently of the considered inverse problem. We show the benefits of the proposed method, in comparison with other regularized DIP methods, for two linear inverse problems, i.e., denoising and super-resolution.**

***Index Terms*—inverse problems, deep image prior, ADMM, denoising, super-resolution**

I. Introduction

Inverse Problems refer to a broad class of problems that can be encountered in many computer vision and image processing problems. The task in inverse problems is to reconstruct a signal from observations that are subject to a known (or inferred) corruption process known as the forward model. Examples of inverse problems are denoising, de-blurring, super-resolution, and reconstruction from a sparse set of measurements. To address inverse problems in 2D imaging, that are typically ill-posed, one has to incorporate some prior knowledge on the kind of typical images we try to restore, which helps restricting the class of admissible solutions.

Early methods for solving inverse problems have been successfully using handcrafted priors such as sparsity or smoothness priors. However, in parallel, task-specific deep neural networks learning the mapping from the measurements to the solution space, have indeed been shown to give significantly higher performances for a variety of applications, e.g., sparse signal recovery [1], deconvolution and deblurring [2], [3], [4], [5], [6], super-resolution [7], [8], [9], and demosaicing [10]. However, it is necessary to learn one network for each type of problems and these networks in general have a very large number of parameters, hence require a very large amount of training data.

The field has hence recently evolved towards coupling classical optimization and deep learning techniques, by using more complex models learned from very large collections of images, instead of simple hand-crafted priors for regularization. In this vein, Chang et al. [11] propose a framework for training a single neural network for solving all linear inverse problems using projected gradient descent. The network serves as a quasi-projection operator for the set of natural images. The projection operator is learned using adversarial learning: a classifier D that fits the decision boundary of the natural image set is used to train a projection function P that maps a signal to the set defined by the classifier. The proposed projection operator is integrated to an Alternating Direction Method of Multipliers (ADMM) algorithm in order to solve arbitrary linear inverse problems. In the plug-and-play ADMM approach [12], the proximal operator of the regularizer is replaced by a traditional denoiser such as BM3D. More recently, the method in [13] uses instead a trained CNN denoiser and solves the problem using the HQS proximal algorithm. A learned proximal operator can also be used to push the estimate in the proximity of a regularized solution. The proximal operator can take the form of a denoising autoencoder [11], [14], [15], or the proximal operator for a regularizer [16]. The authors in [17] learn a generative model of the image space of interest from a collection of training images. The images of interest are assumed to lie on a low-dimensional sub-manifold that can be indexed by an input low dimensional vector.

In parallel, the Deep Image Prior (DIP) has been introduced, as an untrained generative model, to solve inverse problems in image processing, i.e., using no prior information from other images [18, 19]. This approach however may suffer from overfitting to the input measurements, hence the need for early stopping of the optimization, and for regularization. While the original DIP method for solving inverse problems models the set of possible solutions, the authors in [20] instead consider the DIP as a regularizing energy function. Liu et al. [21] introduced a TV constraint regularization when optimizing the DIP to match the observed measurements. Cascarano et al. [22] use the ADMM optimization algorithm to solve the TV-regularized DIP optimization learning problem. While these regularized DIP based solutions avoid having to train a model from a very large collection of images, they under-perform compared with task-specific deep neural networks. To cope with overfitting of the DIP to the input measurements, in particular in the presence of noise, due to the very large number of parameters, the authors in [23] incorporate prior information to regularize the network weights.

In this paper, we study different methods for regularizing the estimate produced by the generative model. While, in the literature, regularization constraints used when optimizing the DIP model are mostly based on handcrafted priors, we consider here a learned regularizer. The goal is to couple advantages of the DIP only optimized from the input image, with those of a generic prior learned from a large collection of images. In that aim, we solve the DIP optimization within the ADMM algorithm in order to use a learned denoiser that plays the role of a proximal operator for regularization. Hence, we refer to this method as DIP-denoiser-ADMM. Experimental results show that for super-resolution, the use of the learned denoiser with ADMM better regularizes the DIP than TV

This work was supported in part by the french ANR research agency in the context of the artificial intelligence project DeepCIM, and in part by the EU H2020 Research and Innovation Programme under grant agreement No 694122 (ERC advanced grant CLIM.

regularization either applied also within ADMM or with a simple gradient descent. The results also show the benefit brought by the generative model (DIP) beyond a direct solution of the inverse problem in an ADMM framework with the same learned proximal operator (i.e. denoiser). For the denoising task, as expected, the direct application of the denoiser remains better than the proposed DIP-denoiser-ADMM. Nevertheless, only a small difference is observed, showing that our method reaches high performance for the two problems of super-resolution and denoising with a unique (hence generic) learned proximal operator.

II. Notations and problem statement

Inverse problems refer to the problems of reconstructing a clean image $x \in \mathbb{R}^d$ from a set of degraded observations, and can be formulated by a linear equation of the form

$$y = Ax + n \tag{1}$$

where $y \in \mathbb{R}^m$ denotes the input observations, $A \in \mathbb{R}^{m \times d}$ represents the degradation operator and $n \in \mathbb{R}^m$ typically represents Additive White Gaussian Noise (AWGN). The degradation matrix A is problem dependent. For example, in the classical problem of image denoising, A is simply the identity matrix. In super-resolution, A represents the blurring and down-sampling operators, whereas in compressive sensing, A is a short-fat matrix with more columns than rows. In general, inverse problems are ill-posed because $rank(A) < d$, that is why we use priors in order to find the optimal solution by solving the optimization problem:

$$\widehat{x} = \underset{x}{\operatorname{argmin}} \ \frac{1}{2} \|y - Ax\|^2 + \lambda \, \phi(x) \tag{2}$$

composed by a data fidelity term and a regularization term representing the chosen signal prior. The former enforces the similarity with the degraded measurements, whereas the latter reflects prior knowledge and a property to be satisfied by the searched solution. The non-negative weighting parameter balances the trade-off between the two terms.

Eq. 2 can be interpreted as the Maximum A Posteriori (MAP) estimation of the clean image x from the observed data y. The MAP maximizes the posterior distribution $p(x|y)$ that can be expressed using the Bayes formula as $p(x|y) = \frac{p(y|x)p(x)}{p(y)}$. The MAP estimate is thus derived as:

$$\begin{aligned} x_{MAP} &= \underset{x}{\operatorname{argmax}} \, p(x|y) = \underset{x}{\operatorname{argmax}} \, p(y|x)p(x), \\ &= \underset{x}{\operatorname{argmin}} \, -log(p(y|x)) - log(p(x)), \end{aligned} \tag{3}$$

where $p(y|x)$ and $p(x)$ are respectively the likelihood and the prior distributions. In the assumption of an Additive White Gaussian Noise n with standard deviation σ, the likelihood is expressed as $p(y|x) = e^{-\frac{\|y-Ax\|_2^2}{2\sigma^2}}$. Hence, by applying the negative logarithm it becomes evident that Eq. 3 is equivalent to Eq. 2 where $\lambda = \sigma^2$ and the regularizer is the negative logarithm of the prior distribution, i.e. $\phi(x) = -log(p(x))$.

Note that for solving Eq. 2, proximal algorithms such as ADMM do not evaluate the regularization term explicitly, but instead require its proximal operator:

$$prox_{\lambda\phi}(u) = \underset{u}{\operatorname{argmin}} \ \frac{1}{2} \|x - u\|^2 + \lambda \, \phi(x) \tag{4}$$

It can be seen that the proximal operator definition is a particular case of inverse problem, where A is simply the identity matrix. Here, the forward model only consists in adding white Gaussian noise. Hence the proximal operator is a Gaussian denoiser in the MAP sense. This property is exploited in plug-and-play methods such as [13, 12] that solve Eq. 2 with a proximal algorithm by replacing the proximal operator with an existing denoiser (e.g. BM3D, trained CNN), instead of defining an explicit regularizer ϕ. The proposed approach also exploits this idea in combination with the DIP method in [18].

III. Regularized Deep Image Prior: Background

While many supervised learning solutions have been proposed for solving inverse problems, we consider here a non supervised method, the deep image prior (DIP) introduced by Ulyanov et al. [18, 19], which can be seen as an image generative model. The deep generative model $x = f_\Theta(z)$ is learned by mapping a random code vector z to an image x, i.e., by solving

$$\Theta^* \in \underset{\Theta}{\operatorname{argmin}} \ \frac{1}{2} \|Af_\Theta(z) - y\|_2^2 \qquad s.t. \ x^* = f_{\Theta^*}(z) \tag{5}$$

where Θ represents the network parameters. The generator is randomly initialized with variables Θ, which are optimized iteratively in a way that the output of the network is close to the degraded measurements. In most of the applications, a U-Net type architecture with skip-connections is used, having over 2 million parameters.

A. TV-regularized DIP

Liu et al. [21] introduced a TV constraint regularization when learning the DIP, which is defined as

$$TV(u) = \sum_{i=1}^{n} \sqrt{(D_h u)_i^2 + (D_v u)_i^2} \tag{6}$$

The variables D_h and D_v denote the first order finite difference discrete operators along the horizontal and vertical axes.
In fact, $\|Af_\Theta(z) - y\|_2^2$ in Eq. 5 can be seen as the data fidelity term in Eq. 2, thus the DIP can be regularized with total variation, yielding the following optimization problem:

$$\begin{aligned} \Theta^* \in \underset{\Theta}{\operatorname{argmin}} \ \frac{1}{2} \|Af_\Theta(z) - y\|_2^2 + \lambda \, TV(f_{\Theta^*}(z)) \\ s.t. \ x^* = f_{\Theta^*}(z) \end{aligned} \tag{7}$$

The objective function is minimized using gradient-based optimization method in [21] and ADMM in [22]. Another difference between these 2 works is that in DIP-TV, anisotropic TV is used, decoupling the contribution of both horizontal and vertical gradient components, whereas ADMM-DIP-TV uses isotropic TV, jointly considering the gradient components.

B. Denoiser-based DIP regularization

Another approach to boost the performance of the DIP is proposed in [24], bringing-in the concept of Regularization by Denoising (RED) [25], which uses existing denoisers for regularizing inverse problems. RED proposes to regularize inverse problems with the following regularization term:

$$\rho(x) = \frac{1}{2} x^T [x - f(x)] \tag{8}$$

where f is the denoiser applied on the candidate image x, and the penalty induced is proportional to the inner-product between this image and its denoising residual.

IV. Regularizing the DIP with a learned denoiser

The proposed method is inspired from methods coupling optimization with learned regularizers [11, 13], but instead, we consider the problem of learning the DIP, regularized by a learned denoiser. Consider the problem of regularizing the DIP optimization with a learned regularizer, therefore minimize:

$$\begin{aligned} \widehat{\Theta} = \underset{\Theta, x}{\operatorname{argmin}} \quad & \frac{1}{2}\|y - A f_\Theta(z)\|_2^2 + \lambda\,\phi(x) \\ \text{subject to} \quad & x = f_\Theta(z) \end{aligned} \tag{9}$$

where ϕ is a learned prior. For simplicity, we will note $t = f_\Theta(z)$ in the following. The augmented Lagrangian form of Eq. 9 can be written as:

$$L(x,t,u) = \frac{1}{2}\|At - y\|_2^2 + \lambda\,\phi(x) + \frac{\rho}{2}\|x - t + u\|_2^2 \tag{10}$$

with ρ being a positive penalty parameter of the constraint, and u the scaled dual variable. This minimization problem is solved using the iterative ADMM method. By alternatively optimizing x, t and u of the augmented Lagrangian $L(x,t,u)$ the ADMM iterate reads as

$$t^{k+1} = f_{\Theta^{k+1}}(z) \quad \text{with}$$

$$\Theta^{k+1} = \underset{\Theta}{\operatorname{argmin}} \|y - A f_\Theta(z)\|_2^2 + \rho \left\|x^k - f_\Theta(z) + u^k\right\|_2^2 \tag{11}$$

$$x^{k+1} = \underset{x}{\operatorname{argmin}} \left\|x - t^{k+1} + u^k\right\|_2^2 + \frac{2\lambda}{\rho}\phi(x) \tag{12}$$

$$u^{k+1} = u^k + x^{k+1} - t^{k+1} \tag{13}$$

As discussed in Sec. II, signal priors are found in the form of proximal operators in ADMM and can be interpreted as Gaussian denoisers. Here, the x-update in Eq. 12 can be rewritten as $x_{k+1} = prox_{\frac{2\lambda}{\rho}\phi}(t^{k+1} - u^k)$. Hence, we propose to replace $prox_{\frac{2\lambda}{\rho}\phi}$ with the learned denoiser from [13]. We thus get

$$x^{k+1} = D_\sigma(t^{k+1} - u^k) \tag{14}$$

with D_σ being the learned denoiser [13] assuming a noise standard deviation $\sigma = \frac{2\lambda}{\rho}$. The denoiser takes the parameter σ as input and was trained using σ values in the range $[0, 50]$. This way, it is applicable for all noise levels in this range while keeping optimal denoising performance for each noise level σ. It is worth noting that since ADMM separates the regularization term ϕ from the degradation matrix A, the learned proximal operator can be used with any linear operator.

In practice, Eq. 11 can't be solved exactly due to the non-linear DIP $f_\Theta(z)$. Furthermore, the DIP should not overfit the exact solution in order to provide an additional regularization effect in complement to the denoiser. Therefore, we solve Eq. 11 inexactly by performing a fixed number n_{DIP} of gradient descent iterations. For the initialization, we optimize the DIP using the original method [18] which performs gradient descent to optimize Eq. 5. A number n_{init} of iterations is used for this initialization, followed by n_{ADMM} iterations of the ADMM scheme.

V. Experiments

In this section, we reproduce the results of [18, 19, 21, 22] on the images of Figure 1 for the problems of denoising and super-resolution and compare them with our proposed method. We also compare the results with the denoiser [13] that we use in our algorithm. In [22], authors use isotropic TV and analyze the addition of ADMM by comparing with anisotropic TV in [21]. For the sake of a fair comparison between [21] and [22], we reproduce the anisotropic version of both of the methods. Table I shows the parameters used for the proposed method for each of the cases of denoising and super-resolution. Figures 2 and 3 show the outputs of denoising with a noise standard deviation $\sigma_0 = 20$ and super-resolution of factor 4 respectively. Tables II and III show the best PSNR along all iterations for each of the compared methods for denoising and super resolution respectively. For denoising (Table II), we can see that, our proposed method significantly improves the result of the DIP. Hence, as expected, learned regularizers outperform handcrafted regularizers when regularizing the DIP optimization. However, the learned denoiser itself performs better alone, and the DIP does not have an added value in this case. This was expected in the case of denoising, since the network is learned end-to-end for the denoising task.

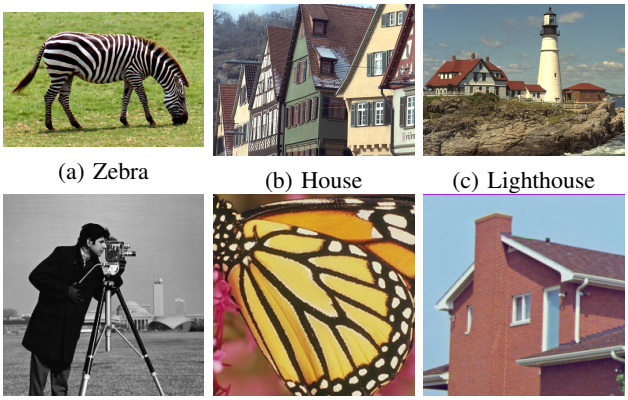

(a) Zebra (b) House (c) Lighthouse

(d) Cameraman (e) Butterfly (f) Hill

Fig. 1: The set of images used for the numerical experiments.

However, for the task of super resolution (Table III), we have a different outcome. In fact, we can see again that regularizing with a learned denoiser over the DIP optimization gives a better performance than using handcrafted priors. But also, as opposed to what we got for the denoising, combining the DIP with the denoiser improves the performance of the learned regularizer alone, which makes our proposed method better than the other ones in this case. It is important to note that the degradation filter for the task of super resolution is a Gaussian filter of standard deviation $\sigma_f = 0.5$ for both x4 and x8 super-resolution factors. More visual results can be seen on the web page: http://clim.inria.fr/DeepCIM/Regul-DIP/index.html

VI. Conclusion

In this paper, we showed the advantages of using a learned denoiser to regularize the Deep Image Prior optimization for inverse problems. The proposed denoiser based regularization compares favourably with handcrafted priors and remain generic, since it can be applied to different inverse problems using an ADMM formulation. We also show that in the context of super-resolution, the use of Deep Image Prior improves on the classical ADMM formulation with the same denoiser based regularization.

References

[1] A. Mousavi and R. G. Baraniuk, "Learning to invert: Signal recovery via deep convolutional networks," in *IEEE int. conf. on acoustics, speech and signal processing (ICASSP)*, pp. 2272–2276, IEEE, 2017.

	(i) Denoising		(ii) Super-Resolution	
	$\sigma_0 = 20$	$\sigma_0 = 30$	Factor 4	Factor 8
n_{init}	1750	1750	2500	4000
n_{ADMM}	100	100	100	100
n_{DIP}	250	250	200	150
σ	20	30	41	45
Step size	0.01	0.01	0.01	0.01

TABLE I: **Parameters used for the proposed method for (i) Denoising with noise standard deviations of ($\sigma_0 = 20$ and $\sigma_0 = 30$), (ii) Super-Resolution factor 4 and 8** . n_{init}: DIP iterations at initialisation, n_{ADMM}: ADMM iterations, n_{DIP}: number of iterations of Eq. 11 in each iteration of ADMM, σ : noise level assumed by the learned denoiser in Eq. 14. The same step size is used for all gradient descent steps.

		DIP	DIPTV	ADMM-DIPTV	denoiser	DIP-denoiser-ADMM
(i) σ_0 = 20	House	27.93	27.81	27.96	31.55	31.35
	Zebra	29.98	29.63	29.5	31.21	31.1
	Lighthouse	29.03	29.55	29.3	32.18	31.96
	Hill	31.29	32.53	32.07	34.02	33.93
	Butterfly	30.07	30.52	30.62	33.23	32.96
	Cameraman	30.01	31.4	31.04	33.7	33.67
	Average	29.72	30.24	30.08	32.64	32.49
(ii) σ_0 = 30	House	25.75	26.18	25.72	29.28	29.15
	Zebra	27.34	27.52	27.58	28.93	28.87
	Lighthouse	26.7	27.48	27.2	29.96	29.75
	Hill	29.61	30.81	30.28	32.62	32.51
	Butterfly	28.16	28.48	28.52	30.85	30.83
	Cameraman	27.79	28.88	28.41	31.15	31.14
	Average	27.55	28.23	27.95	30.46	30.37

TABLE II: **PSNR of denoised images obtained with TV, DIP [18], DIPTV [21] , ADMM-DIPTV [22], the denoiser of [13] and our proposed method, for noise standard deviation (i) 20 and (ii) 30.**

		Bicubic	DIP	DIPTV	ADMM-DIPTV	denoiser-ADMM	DIP-denoiser-ADMM
(i) factor 4	House	19.42	19.72	19.85	19.62	20.06	20.14
	Zebra	23.00	24.13	24.31	23.84	24.71	24.9
	Lighthouse	23.09	23.05	23.37	23.17	23.46	23.45
	Hill	26.37	28.44	27.61	27.01	28.75	28.97
	Butterfly	20.95	24.23	23.66	22.47	24.34	24.42
	Cameraman	22.58	23.23	23.27	22.99	23.34	23.51
	Average	22.58	23.8	23.67	23.18	24.11	24.23
(ii) factor 8	House	16.25	17.16	17.32	17.11	16.83	17.21
	Zebra	16.84	18.01	18.15	17.75	17.37	17.84
	Lighthouse	19.87	20.6	20.65	20.45	20.47	20.57
	Hill	20.75	22.45	22.1	22.01	23.46	23.58
	Butterfly	15.38	17.04	16.62	16.63	17.11	17.4
	Cameraman	18.94	20.2	20.1	19.95	19.88	20.19
	Average	18.00	19.24	19.15	18.98	19.18	19.46

TABLE III: **PSNR of upsampled images obtained with bicubic interpolation, DIP [18], DIPTV [21] , ADMM-DIPTV [22], the denoiser of [13] with ADMM and our proposed method, for (i) SR factor 4 and (ii) SR factor 8.**

[2] L. Xu, J. S. Ren, C. Liu, and J. Jia, "Deep convolutional neural network for image deconvolution," *Advances in neural information processing systems*, vol. 27, pp. 1790–1798, 2014.

[3] Y. Tan, D. Zhang, F. Xu, and D. Zhang, "Motion deblurring based on convolutional neural network," in *Int. Conf. on Bio-Inspired Computing: Theories and Applications. Springer*, pp. 623–635, 2017.

[4] K. Uchida, M. Tanaka, and M. Okutomi, "Non-blind image restoration based on convolutional neural network," in *IEEE Global Conf. on Consumer Electronics (GCCE)*, pp. 40–44, IEEE, 2018.

[5] T. Eboli, J. Sun, and J. Ponce, "End-to-end interpretable learning of non-blind image deblurring," *arXiv preprint arXiv:2007.01769*, 2020.

[6] J. Dong, S. Roth, and B. Schiele, "Deep wiener deconvolution: Wiener meets deep learning for image deblurring," *Advances in Neural Information Processing Systems*, vol. 33, 2020.

[7] C. Dong, C. C. Loy, K. He, and X. Tang, "Learning a deep convolutional network for image super-resolution," in *European conf. on computer vision*, pp. 184–199, Springer, 2014.

[8] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. P. Aitken, A. Tejani, J. Totz, and Z. Wang, "Photo-realistic single image super-resolution using a generative adversarial network," in *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, p. 4681–4690, 2017.

[9] Q. Ning, W. Dong, G. Shi, L. Li, and X. Li, "Accurate and lightweight image super-resolution with model-guided deep unfolding network," *IEEE Journal of Selected Topics in Signal Processing*, 2020.

[10] B. Henz, E. S. Gastal, and M. M. Oliveira, "Deep

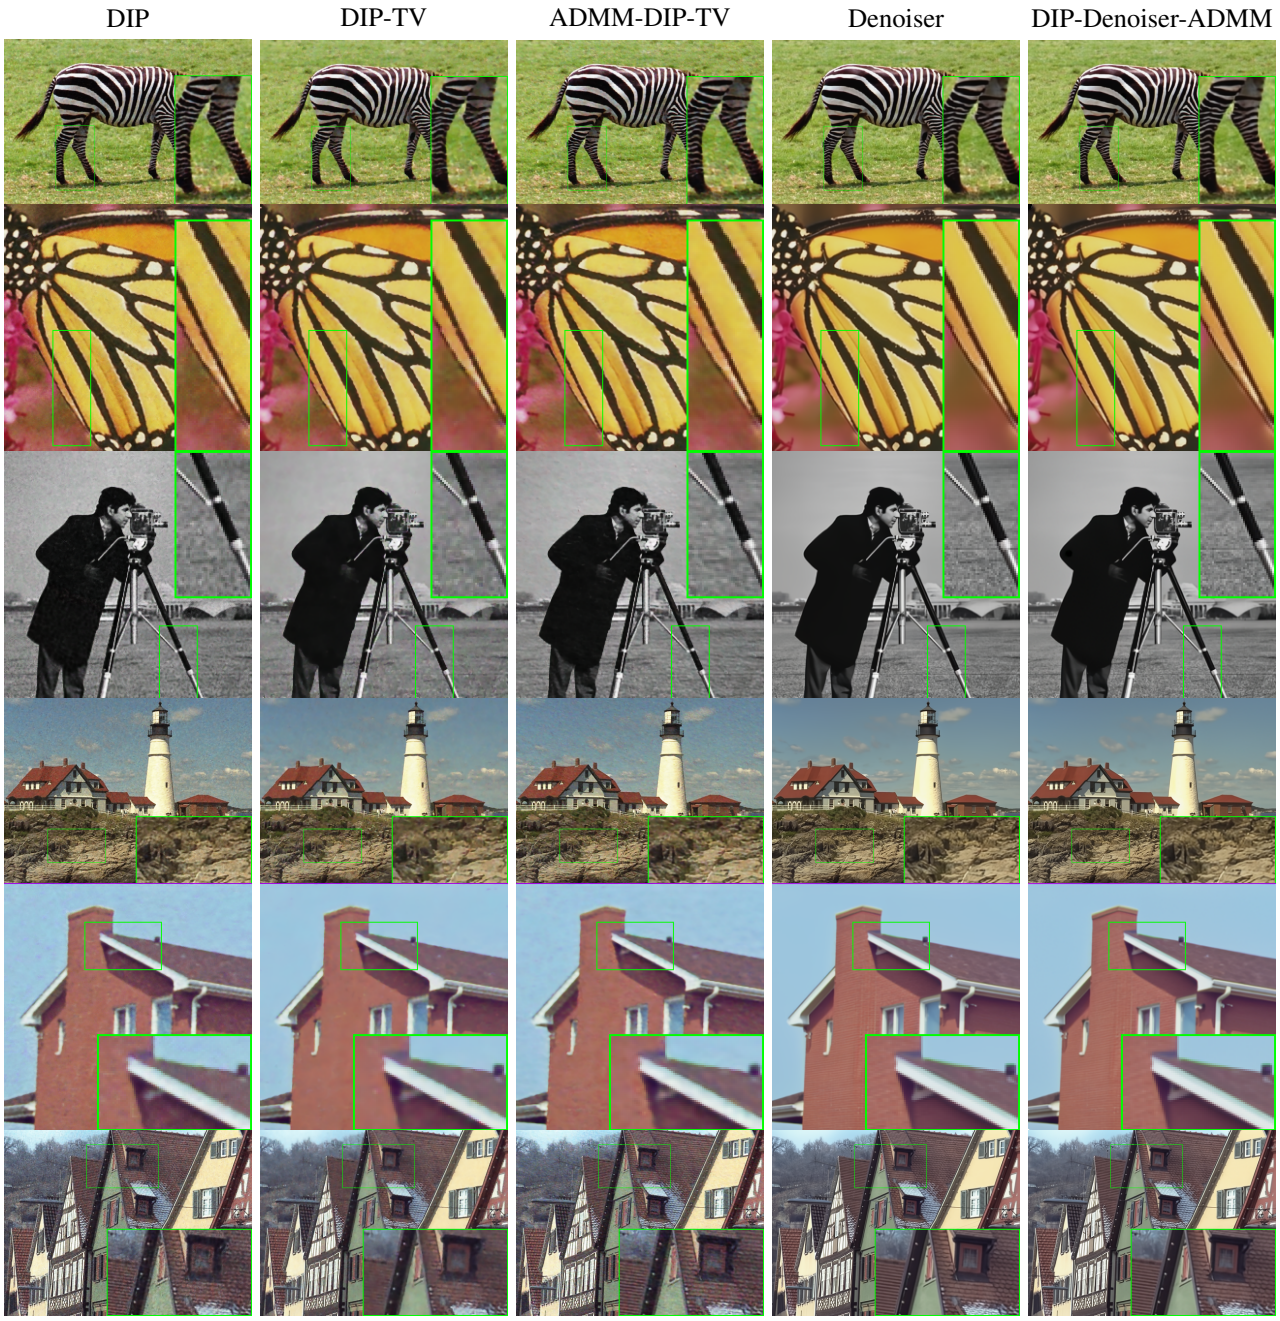

Fig. 2: Results obtained for the denoising application with a noise of standard deviation $\sigma_0 = 20$

joint design of color filter arrays and demosaicing," in *Computer Graphics Forum*, vol. 37, pp. 389–399, Wiley Online Library, 2018.

[11] J. H. R. Chang, C.-L. Li, B. Poczos, B. V. K. V. Kumar, and A. C. Sankaranarayanan, "One network to solve them all — solving linear inverse problems using deep projection models," *IEEE Int. Conf. on Computer Vision (ICCV)*, 2017.

[12] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *IEEE Global Conf. on Signal and Information Processing*, pp. 945–948, 2013.

[13] K. Zhang, Y. Li, W. Zuo, L. Zhang, L. Van Gool, and R. Timofte, "Plug and-play image restoration with deep denoiser prior," 2020.

[14] H. Gupta, K. H. Jin, H. Q. Nguyen, M. T. McCann, and M. Unser, "Cnn-based projected gradient descent for consistent ct image reconstruction," *IEEE Trans. on Medical Imaging*, vol. 37, no. 6, p. 1440–1453, 2018.

[15] T. Meinhardt, M. Moeller, C. Hazirbas, and D. Cremers, "Learning proximal operators: Using denoising networks for regularizing inverse imaging problems," in *IEEE Int. Conf. on Computer Vision (ICCV)*, pp. 1799–1808, 2017.

[16] Y. Chen and T. Pock, "Trainable nonlinear reaction diffusion: A flexible framework for fast and effective image restoration," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 39, no. 6, p. 1256–1272, 2017.

[17] A. Bora, A. Jalal, E. Price, and A. G. Dimakis, "Compressed sensing using generative models," in *Int. Conf. on Machine Learning (ICML)*, pp. 537–546, 2017.

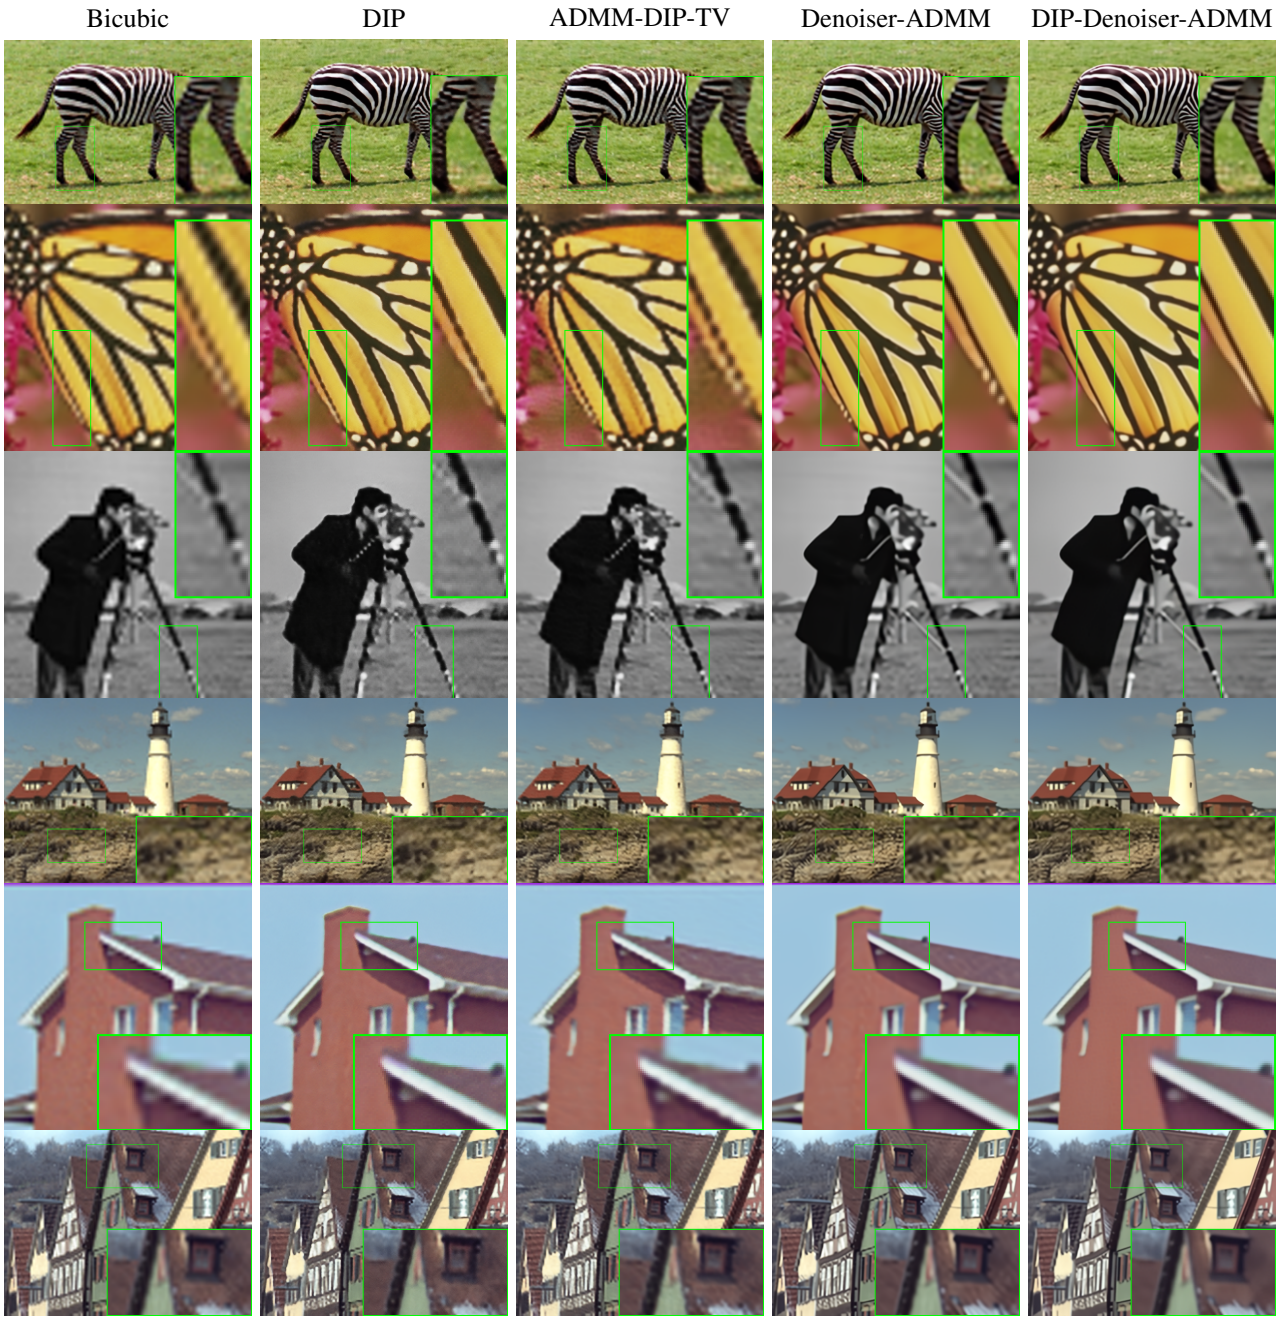

Fig. 3: Results obtained for the super-resolution application with magnifying factor of 4. The low resolution input image is generated using a Gaussian filtering kernel of standard deviation $\sigma_f = 0.5$ followed by a downsampling of 4.

[18] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Deep image prior," *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[19] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Deep image prior," *Int. Journal of Computer Vision (IJCV)*, 2018.

[20] A. Sagel, A. Roumy, and C. Guillemot, "Sub-dip: Optimization on a subspace with deep image prior regularization and application to superresolution," *IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2513–2517, 2020.

[21] J. Liu, Y. Sun, X. Xu, and U. S. Kamilov, "Image restoration using total variation regularized deep image prior," *IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP) pages = 7715–7719*, 2019.

[22] P. Cascarano, A. Sebastiani, M. C. Comes, G. Franchini, and F. Porta, "Combining weighted total variation and deep image prior for natural and medical image restoration via admm," *arXiv preprint arXiv:2009.11380*, 2021.

[23] D. V. Veen, A. Jalal, E. Price, S. Vishwanath, and A. Dimakis, "Compressed sensing with deep image prior and learned regularization," *ArXiv*, vol. abs/1806.06438, 2018.

[24] G. Mataev, M. Elady, and P. Milanfarz, "Deepred: Deep image prior powered by red," *arXiv preprint arXiv:1903.10176*, 2019.

[25] Y. Romano, M. Elad, and P. Milanfar, "The little engine that could regularization by denoising (red)," *SIAM Journal on Imaging Sciences*, vol. 10, pp. 1804–1844, 2017.