Constrained Regularization by Denoising with Automatic Parameter Selection

Pasquale Cascarano,* Alessandro Benfenati,† Ulugbek S. Kamilov,‡ *Senior Member, IEEE*, Xiaojian Xu§

January 15, 2024

Abstract

Regularization by Denoising (RED) is a well-known method for solving image restoration problems by using learned image denoisers as priors. Since the regularization parameter in the traditional RED does not have any physical interpretation, it does not provide an approach for automatic parameter selection. This letter addresses this issue by introducing the Constrained Regularization by Denoising (CRED) method that reformulates RED as a constrained optimization problem where the regularization parameter corresponds directly to the amount of noise in the measurements. The solution to the constrained problem is solved by designing an efficient method based on alternating direction method of multipliers (ADMM). Our experiments show that CRED outperforms the competing methods in terms of stability and robustness, while also achieving competitive performances in terms of image quality.

Keywords: Image restoration, plug-and-play priors, regularization by denoising, discrepancy principle.

1 Introduction

The problem of recovering an image $\mathbf{x} \in \mathbb{R}^n$ from its degraded measurement $\mathbf{b} \in \mathbb{R}^n$ can be cast as the following linear inverse problem

$$\text{find} \quad \mathbf{x} \in \mathbb{R}^n \quad \text{such that} \quad \mathbf{b} = \mathbf{A}\mathbf{x} + \boldsymbol{\eta}, \tag{1}$$

where $\mathbf{A} \in \mathbb{R}^{n \times n}$ is a known measurement operator and $\boldsymbol{\eta} \in \mathbb{R}^n$ is random noise with standard deviation $\sigma_{\boldsymbol{\eta}}$.

Linear inverse problems are at the core of many applications [1–5]. However, since most inverse problems are ill-posed, it is common to formulate the solution $\mathbf{x}^* \in \mathbb{R}^n$ of (1) as a minimizer of a regularized objective function

$$\mathbf{x}^* = \operatorname*{argmin}_{\mathbf{x} \in \mathbb{R}^n} \ell(\mathbf{x}; \mathbf{b}) + \lambda \rho(\mathbf{x}). \tag{2}$$

The data-fidelity term $\ell(\mathbf{x}; \mathbf{b})$ encodes information on the noise statistics, e.g. additive white Gaussian noise (AWGN) assumptions entail $\ell(\mathbf{x}; \mathbf{b}) = \frac{1}{2}\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2$. The regularization parameter $\lambda > 0$ is often hand-tuned to obtain optimal restored images. Alternatively, it can be estimated using well established methods such as the discrepancy principle, L-curve, or cross-validation [6]. However, the primary challenge consists in designing the regularization functional $\rho(\mathbf{x})$ in order to capture the intricate image features.

Plug-and-Play (PnP) Priors framework has recently emerged as a powerful tool for exploiting sophisticated denoisers as regularizers without explicitly defining $\rho(\mathbf{x})$ [7–13]. However, the lack of an explicit objective function complicates the analysis of PnP methods in terms of theoretical understanding and convergence guarantees. Regularization by denoising (RED) [14] is a variant of PnP based on formulating an explicit regularization functional

$$\rho_{\text{RED}}(\mathbf{x}) := \frac{1}{2} \left(\mathbf{x}^T \left(\mathbf{x} - f(\mathbf{x}) \right) \right), \tag{3}$$

where $f(\cdot)$ denotes a denoiser. When the denoiser is differentiable, locally homogeneous, and has a symmetric Jacobian [14, 15], ρ_{RED} is convex and its gradient can be efficiently computed as $\nabla \rho_{\text{RED}}(\mathbf{x}) = \mathbf{x} - f(\mathbf{x})$. The interpretation with an explicit regularizer simplifies the theoretical analysis of RED algorithms as methods for computing global minimizers of convex objective functions. These conditions are often not satisfied for many practical denoisers [15], however, the RED algorithms achieve state-of-the-art performances in many imaging applications.

RED was recently reformulated as a constrained optimization problem by projecting the least square minimum onto the fixed-point sets of demicontractive denoisers, which are proven to be convex sets [16]. Ideally, an image denoiser f should satisfy the condition $\mathcal{M} \subset \text{Fix}(f)$, where $\mathcal{M}$ denotes the manifold of natural images. In practice, denoisers are often far from being ideal, and their fixed-point sets may not correspond to the set of natural images, leading to suboptimal recovery results. Additionally, even if any non-expansive denoiser is demicontractive, determining whether a given denoiser is non-expansive or demicontractive is a challenge.

In this letter, we present a constrained RED (CRED) approach inspired by the discrepancy principle [17, 18] solved via the alternating direction method of multipliers (ADMM). In order to overcome the limits regarding $\text{Fix}(f)$, we reverse the RED-PRO [16] formulation by considering the RED regularization functional subject to the discrepancy between the measured data and the reconstruction being below a given threshold. This threshold represents the strength of the regularization and has a precise physical significance since it reflects the standard deviation of the noise affecting the data. Different from the regularization parameter of the original unconstrained RED formulation [14], which must be hand-tuned, the threshold can be estimated by using well-established techniques [19]. Therefore, our approach avoids any parameter tuning which may be limiting in real applications. Our formulation and the corresponding ADMM scheme are presented in Section 2. In Section 3, we underline the quality of CRED in terms of image quality metrics and robusteness, with respect to the choice of the denoiser and the ADMM parameters, through several comparisons with RED and RED-PRO.

2 Constrained RED

Upon AWGN assumptions, CRED seeks to compute a solution of (1) by solving the constrained problem

$$\underset{\mathbf{x} \in \mathbb{R}^n}{\text{argmin}} \ \rho_{\text{RED}}(\mathbf{x}) \quad \text{subject to} \quad \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 \leq \delta, \tag{4}$$

where $\delta := \tau \sqrt{n} \, \sigma_{\boldsymbol{\eta}}$ with $\tau \in [0, 1]$, n number of pixels in the image and $\sigma_{\boldsymbol{\eta}}$ is the noise level, which is assumed to be known. Since $\rho_{\text{RED}}(\cdot)$ is continuous and the constraint set is also bounded, the problem (4) has at least one solution by the Weierstrass theorem. The problem (4) can be equivalently reformulated into the following form

$$\underset{\mathbf{x},\mathbf{t},\mathbf{r} \in \mathbb{R}^n}{\text{argmin}} \ \rho_{\text{RED}}(\mathbf{t}) + \iota_{B_\delta}(\mathbf{r}), \ \text{s.t} \ \mathbf{r} = \mathbf{A}\mathbf{x} - \mathbf{b}, \mathbf{x} = \mathbf{t}, \tag{5}$$

where $B_\delta := \{\mathbf{r} \in \mathbb{R}^n \mid \|\mathbf{r}\|_2^2 \leq \delta\}$ and ι_{B_δ} denotes the indicator function of the set B_δ. The corresponding Augmented Lagrangian function is given by

$$\begin{aligned}
L(\mathbf{x}, \mathbf{r}, \mathbf{t}; \boldsymbol{\lambda_t}, \boldsymbol{\lambda_r}) = {} & \rho_{\text{RED}}(\mathbf{t}) + i_{B_\delta}(\mathbf{r}) \\
& + \frac{\beta_{\mathbf{r}}}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b} - \mathbf{r} + \frac{\boldsymbol{\lambda_r}}{\beta_{\mathbf{r}}}\|_2^2 \\
& + \frac{\beta_{\mathbf{t}}}{2} \left\| \mathbf{x} - \mathbf{t} + \frac{\boldsymbol{\lambda_t}}{\beta_{\mathbf{t}}} \right\|_2^2 - \frac{1}{2}\|\boldsymbol{\lambda_r}\|_2^2 - \frac{1}{2}\|\boldsymbol{\lambda_t}\|_2^2,
\end{aligned} \tag{6}$$

2

where $\boldsymbol{\lambda}_{\mathbf{r}}$ and $\boldsymbol{\lambda}_{\mathbf{t}}$ are the Lagrange multipliers, and $\beta_{\mathbf{r}}$ and $\beta_{\mathbf{t}}$ are real positive penalties. The optimization problem involving the Augmented Lagrangian can then solved via the Alternating Direction Method of Multipliers (ADMM) [20–22]. Algorithm 1 summarizes the ADMM method.

Algorithm 1 Constrained RED Approach (CRED)

1: Set δ, $\mathbf{x}^0$, $\mathbf{t}^0 = 0$ and $\mathbf{r}^0 = \mathbf{b} - \mathbf{A}\mathbf{x}^0$, select $\beta_{\mathbf{u}}, \beta_{\mathbf{r}} > 0$ and initialise $\boldsymbol{\lambda}_{\mathbf{r}}^0, \lambda_{\mathbf{t}}^0$
2: **for** $k = 0, 1, \ldots$ **do**

3: $\quad \mathbf{x}^{k+1/2} = \dfrac{\beta_{\mathbf{r}}}{\beta_{\mathbf{t}}} \mathbf{A}^\top \left(\mathbf{b} + \mathbf{r}^k - \dfrac{\boldsymbol{\lambda}_{\mathbf{r}}^k}{\beta_{\mathbf{r}}} \right) + \left(\mathbf{t}^k - \dfrac{\boldsymbol{\lambda}_{\mathbf{t}}^k}{\beta_{\mathbf{t}}} \right)$

4: $\quad \mathbf{x}^{k+1} \leftarrow \left(\dfrac{\beta_{\mathbf{r}}}{\beta_{\mathbf{t}}} \mathbf{A}^\top \mathbf{A} + \mathbf{I} \right)^{-1} \mathbf{x}^{k+1/2}$

5: $\quad \mathbf{t}^{k+1} = \dfrac{1}{1+\beta_{\mathbf{t}}} f(\mathbf{t}^k) + \dfrac{\beta_{\mathbf{t}}}{1+\beta_{\mathbf{t}}} \left(\mathbf{x}^{k+1} + \dfrac{\boldsymbol{\lambda}_{\mathbf{t}}^k}{\beta_{\mathbf{t}}} \right)$

6: $\quad \mathbf{r}^{k+1} = \mathrm{proj}_{\mathcal{B}_\delta} \left(\mathbf{A}\mathbf{x}^{k+1} - \mathbf{b} + \dfrac{\boldsymbol{\lambda}_{\mathbf{r}}^k}{\beta_{\mathbf{r}}} \right)$

7: $\quad \boldsymbol{\lambda}_{\mathbf{r}}^{k+1} \leftarrow \boldsymbol{\lambda}_{\mathbf{r}}^k + \beta_{\mathbf{r}}(\mathbf{r}^{k+1} - \mathbf{A}\mathbf{x}^{k+1} + \mathbf{b})$

8: $\quad \boldsymbol{\lambda}_{\mathbf{t}}^{k+1} \leftarrow \boldsymbol{\lambda}_{\mathbf{t}}^k + \beta_{\mathbf{t}} \left(-\mathbf{x}^{k+1} + \mathbf{t}^{k+1} \right)$
9: **end for**

Remark 1. *It is worth highlighting the following points relative to Algorithm 1. In Line 1 the definition of δ requires the knowledge of $\sigma_{\boldsymbol{\eta}}$. This is not limiting since a good estimation of the noise standard deviation can often be obtained [19]. In lines 3-4 the optimality conditions on the $\ell_2 - \ell_2$ subproblem are used. This subproblem can often be solved using FFT by imposing periodic boundary conditions. In line 5: we adopt the fixed-point strategy, by zeroing the derivative as in [23]. In line 6: $\mathrm{proj}_{\mathcal{B}_\delta}$ refers to the projection onto $\mathcal{B}_\delta$. In the implementation, we adopt the variable change $\lambda \sim \frac{\lambda}{\beta}$ for simplifying the notation. Moreover, due to the convexity of ρ_{RED}, the algorithm converges to the minimum of (4).*

3 Numerical Results

3.1 Settings, evaluation metrics, and baseline methods

We focus on the task of image deblurring with AWGN. Therefore, $\mathbf{A}$ in (1) is a Gaussian blurring operator of standard deviation $\sigma_{\mathbf{A}}$.

We simulate blurry and noisy data by applying the image formation model (1) to the images from Set5 [24] and Set24 [25] referred to as ground truths (GTs). We compare our method with two baselines: the original RED formulation [14] solved using ADMM and its more recent variant RED-PRO [16] solved via gradient-descent. We investigate their behaviour with respect to the choice of some hyperparameters: we focus on the role of λ and (α, μ) for RED and RED-PRO, respectively. The former is the regularization parameter, the couple (α, μ) represent the strength of the denoiser and the starting steplength, respectively. Concerning CRED, we consider penalties such that $\beta_{\mathbf{r}}^{k+1} = \gamma \beta_{\mathbf{r}}^k$ and $\beta_{\mathbf{t}}^{k+1} = \gamma \beta_{\mathbf{t}}^k$. In the following, we comment about the choice of γ and δ.

For all the methods, we consider the relative difference of the iterates with tolerance equals 10^{-4} as the stopping criterion. We set the maximum number of iterations to 200.

We consider the set of CNN based denoisers introduced in [26]. In order to evaluate the influence of the strength of the denoiser, we choose the ones trained for removing Gaussian noise of standard deviation equal to 16 and 30, which are referred to as $\mathbf{D}_1$ and $\mathbf{D}_2$, respectively.

To assess the quality of the restored images we consider the PSNR and SSIM metrics. Moreover, we point out that from a theoretical perspective, given a ground truth image $\mathbf{x}$ and its blurred and noisy simulated data $\mathbf{b}$, $\sigma_{\mathbf{x}} := \frac{\|\mathbf{A}\mathbf{x}-\mathbf{b}\|_2}{\sqrt{n-1}}$ represents the unbiased estimator of $\sigma_{\boldsymbol{\eta}}$. For this reason we consider as valuable metric the comparison between the real noise standard deviation $\sigma_{\boldsymbol{\eta}}$ and $\sigma_{\mathbf{x}^*}$ where $\mathbf{x}^*$ refers to the output of the algorithms.

Table 1: Mean relative errors between $\sigma_{\mathbf{x}^*}$ and $\sigma_{\boldsymbol{\eta}}$ for all the images in Set5.

Metric	$\sigma_{\boldsymbol{\eta}} = 15$	$\sigma_{\boldsymbol{\eta}} = 25$	$\sigma_{\boldsymbol{\eta}} = 35$	$\sigma_{\boldsymbol{\eta}} = 50$
$\mathrm{RE}(\sigma_{\boldsymbol{\eta}}, \sigma_{\mathbf{x}^*})$	0.0050	0.0045	0.0044	0.0043

3.2   On the choice of the threshold

In this section, we consider Set5, set $\sigma_{\mathbf{A}} = 1$, $\sigma_{\boldsymbol{\eta}} = 25$, $\gamma = 1.01$, and use only $\mathbf{D}_1$ as the denoiser. We inspect the influence of $\delta = \tau \sqrt{n} \, \sigma_{\boldsymbol{\eta}}$, with $\tau \in [0, 1]$ on the restored images. We consider two different scenarios: in the first we assume $\sigma_{\boldsymbol{\eta}}$ is known (idealized scenario), whereas in the second we assume only an estimate $\overline{\sigma_{\boldsymbol{\eta}}}$ is provided (realistic scenario). The latter is computed as in [19].

(a) Plot of $\sigma_{x^*}(\tau)$.   (b) Plot of PSNR(τ).
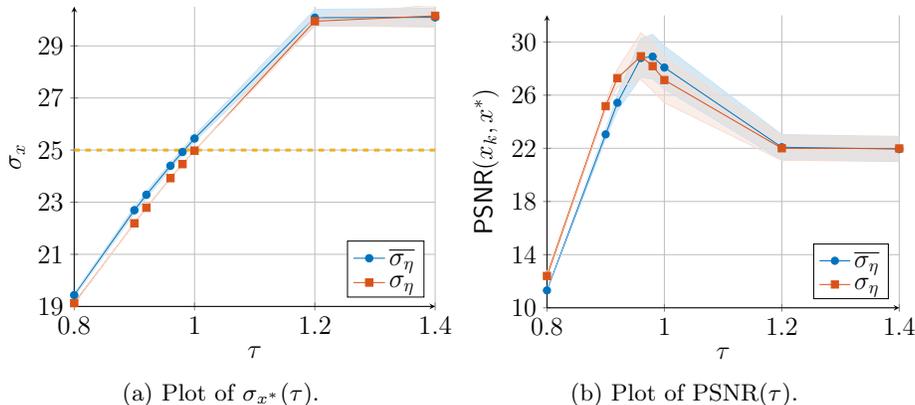
Figure 1: Distribution of $\sigma_{\mathbf{x}^*}$ (a) and PSNR (b) by varying τ for idealized and realistic scenarios on the whole Set5.

In Figure 1(a) and Figure 1(b) the continuous lines show the mean of the distribution of $\sigma_{\mathbf{x}^*}$ and PSNR as function of τ. Moreover, we shade the region spanned by the standard deviation of their distributions. The orange and blue lines represent the idealized and realistic scenarios, respectively. The yellow dashed line represents the noise standard deviation $\sigma_{\boldsymbol{\eta}}$ affecting the simulated data.

As expected, in the idealized scenario, we can observe that $\sigma_{\mathbf{x}^*}$ approximates $\sigma_{\boldsymbol{\eta}}$ when $\tau = 1$ for all images in Set5. Conversely, when only an estimate $\overline{\sigma_{\boldsymbol{\eta}}}$ is given we observe that the best approximation of $\sigma_{\boldsymbol{\eta}}$ is reached when $\tau = 0.98$. We hypothesize that this is due to the fact that the algorithm in [19] tends to overestimate the noise level in the simulated data $\mathbf{b}$. Moreover, Figure 1(a) shows that small values of τ underestimate, whereas large values of τ overestimate $\sigma_{\boldsymbol{\eta}}$. In Figure 1(b) we show the behaviour of the PSNR metric by varying τ. We obtain comparable performances in terms of PSNR for both scenarios.

For all the following experiments we set $\tau = 0.98$ and we estimate $\sigma_{\boldsymbol{\eta}}$ by [19]. In Table 1 we consider different level of degradations ($\sigma_{\mathbf{A}} = 1, \sigma_{\boldsymbol{\eta}} = 15, 25, 35, 50$) and we report the mean of the relative errors (RE) between $\sigma_{\mathbf{x}^*}$ and $\sigma_{\boldsymbol{\eta}}$ for all the images in Set5. We observe that the mean of the relative errors is small (less than 0.5%) while changing the level of degradation.

3.3   On the choice of the ADMM penalties and denoiser

In this section we consider the sole *Butterfly* image by setting $\sigma_{\mathbf{A}} = 1.2$ and $\sigma_{\boldsymbol{\eta}} = 30$. We investigate the stability of CRED with respect to the choice of γ, $\beta_{\mathbf{r}}^0$ and $\beta_{\mathbf{t}}^0$. We point out that in the experiments we did not observe any significant difference when choosing different images.

In Table 2 we report the mean values of the distribution of PSNR and SSIM obtained by choosing $\left(\beta_{\mathbf{r}}^0, \beta_{\mathbf{t}}^0\right) \in \{0.2, 0.4, 0.6, 0.8, 1\}^2$ for different γ. In this experiment we consider $\mathbf{D}_1$ as denoiser. As a general comment, we observe that we can reach similar performances in terms of PSNR and SSIM for the considered values of γ.

Table 2: Mean values of PSNR and SSIM distributions over different configurations of $\left(\beta_{\mathbf{r}}^0, \beta_{\mathbf{t}}^0\right)$ for different γ.

Metric	$\gamma = 1$	$\gamma = 1.01$	$\gamma = 1.05$
PSNR	24.6605	24.7312	24.4178
SSIM	0.9342	0.9324	0.9219

(a) RED	(b) REDPRO	(c) CRED
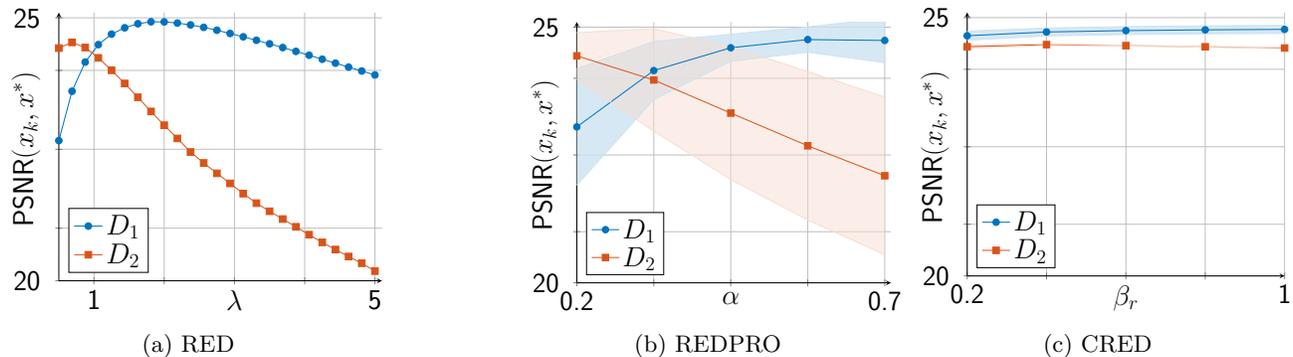
Figure 2: PSNR behaviour by varying λ for RED (a), (α, μ) and μ for RED-PRO (b), $(\beta_{\mathbf{r}}, \beta_{\mathbf{t}})$ for CRED (c). In (b) for each α we present the PSNR distribution (average $\pm$ std) wrt μ; in (c) for each $\beta_{\mathbf{r}}$ we present the PSNR distribution (average $\pm$ std) wrt $\beta_{\mathbf{t}}$.

Figure 2(c) depicts the distribution of the PSNR values when setting $\gamma = 1.01$ for the CNN denoisers $\mathbf{D}_1$ (orange) and $\mathbf{D}_2$ (blue). CRED appears stable regardless the considered choices of $\beta_{\mathbf{r}}^0$ and $\beta_{\mathbf{t}}^0$, and moreover, it seems that the selection of CNN denoisers has a minimal impact on the overall performances. In the following experiments we fix $\beta_{\mathbf{r}}^0 = 1$, $\beta_{\mathbf{t}}^0 = 1$ and $\gamma = 1.01$.

3.4   Comparisons with RED and RED-PRO

In this section, we compare CRED, RED and RED-PRO in terms of stability with respect to their hyperparameters and reconstruction metric on the Set5 and Set24.

In order to compare the stability of RED and RED-PRO with respect to CRED, we consider the simulated *Butterfly* image used in the Subsection 3.3. For the original RED algorithm, we sample 25 different values of $\lambda \in [0.5, 5]$. Concerning RED-PRO, we consider these 25 different configurations $(\alpha, \mu) \in \{0.3, 0.4, 0.5, 0.6, 0.7\}^2$.

In Figure 2(a) we plot the PSNR behaviour of RED by varying λ, whereas in Figure 2(b) we report the PSNR distribution of RED-PRO for different configurations (α, μ). The blue lines represent the case where the denoiser $\mathbf{D}_1$ is used, whereas the orange lines represent the case where the denoiser $\mathbf{D}_2$ is used.

By comparing these results with the ones reported in Figure 2(c), it is evident how our CRED looks more stable with respect to the choice of his hyperparameters. Finally, we observe that for RED and RED-PRO the configuration of parameters maximising the PSNR changes when considering different denoisers. We stress that the same conclusions apply when considering different images.

In order to compare the reconstruction metric performances we consider all the images from Set5 and Set24 for different degradation levels. In Table 3 we report the mean PSNR and SSIM values. For the competing methods RED and RED-PRO the hyperparameters have been estimated in order to minimize the difference between $\sigma_{\boldsymbol{\eta}}$ and $\sigma_{\mathbf{x}^*}$. We notice that in terms of the considered metrics CRED performs as well as RED and RED-PRO. However, we remark that it does not require any parameter tuning.

In Figure 3(a) and Figure 3(b) we report the GT and the degraded close-ups of *Baby* from Set5. Moreover, in Figures 3(c), 3(d) and 3(e) we depict the restored close-ups by RED, RED-PRO and CRED. In terms of visual quality there are no relevant differences between the restored images however CRED seems to slightly reproduce more clearly image details than RED and RED-PRO.

Table 3: Mean values of PSNR and SSIM for the images in Set5 and Set24 by varying the degradation levels.

Metric	Set5 ($\sigma_{\mathbf{A}} = 1, \sigma_\eta = 15$)			Set24 ($\sigma_{\mathbf{A}} = 1.2, \sigma_\eta = 25$)		
	RED	PRORED	CRED	RED	PRORED	CRED
PSNR	30.04	29.61	30.05	26.70	26.29	26.95
SSIM	0.91	0.90	0.91	0.77	0.76	0.77

GT   Degraded   RED   RED-PRO   CRED
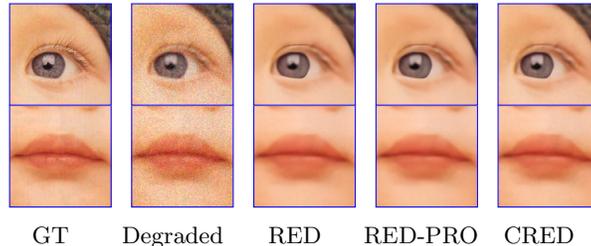
Figure 3: Restoration of the *Baby* image. From left to right: two close-ups of ground truth, degraded image, RED, RED-PRO, and CRED.

4 Conclusions

This letter presents a novel constrained formulation of the popular RED method that forces the minimum of the regularization functional to satisfy a discrepancy-based bound for a given threshold serving as regularization parameter. Our formulation is then solved within the ADMM framework and the overall approach results in a simple yet effective method for image restoration, which is called CRED through the letter. Defining the threshold requires an estimate of the standard deviation of the noise affecting the data which can be provided as described in [19] thus eliminating the need for extensive parameter estimation. The key point of CRED is its superior stability and robustness with respect to both model and algorithm hyperparameters which is assessed through several comparisons with the original RED and its variant RED-PRO methods. Furthermore, the experiments conducted show that in terms of PSNR and SSIM metrics CRED performs as well as, if not better, when compared to both original RED and RED-PRO. Finally, performances, stability and robustness make CRED a promising choice for image restoration.

References

[1] P. Cascarano, F. Corsini, S. Gandolfi, E. L. Piccolomini, E. Mandanici, L. Tavasci, and F. Zama, "Super-resolution of thermal images using an automatic total variation based method," *Remote Sensing*, vol. 12, no. 10, p. 1642, 2020.

[2] P. Cascarano, M. C. Comes, A. Sebastiani, A. Mencattini, E. Loli Piccolomini, and E. Martinelli, "Deepcel0 for 2D single-molecule localization in fluorescence microscopy," *Bioinformatics*, vol. 38, no. 5, pp. 1411–1419, 2022.

[3] A. M. Teodoro, J. M. Bioucas-Dias, and M. A. Figueiredo, "A convergent image fusion algorithm using scene-adapted gaussian-mixture-based denoising," *IEEE Transactions on Image Processing*, vol. 28, no. 1, pp. 451–463, 2018.

[4] D. Mylonopoulos, P. Cascarano, L. Calatroni, and E. L. Piccolomini, "Constrained and unconstrained inverse Potts modelling for joint image super-resolution and segmentation," *Image Processing On Line*, vol. 12, pp. 92–110, 2022.

[5] A. Benfenati and V. Ruggiero, "Image regularization for Poisson data," *Journal of Physics: Conference Series*, vol. 657, no. 1, p. 012011, nov 2015.

[6] M. Bertero, "Regularization methods for linear inverse problems," in *Inverse Problems: Lectures given at the 1st 1986 Session of the Centro Internazionale Matematico Estivo (CIME) held at Montecatini Terme, Italy, May 28–June 5, 1986.* Springer, 2006, pp. 52–112.

[7] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *2013 IEEE global conference on signal and information processing.* IEEE, 2013, pp. 945–948.

[8] P. Cascarano, E. L. Piccolomini, E. Morotti, and A. Sebastiani, "Plug-and-play gradient-based denoisers applied to CT image enhancement," *Applied Mathematics and Computation*, vol. 422, p. 126967, 2022.

[9] Y. Sun, Z. Wu, X. Xu, B. Wohlberg, and U. S. Kamilov, "Scalable plug-and-play ADMM with convergence guarantees," *IEEE Transactions on Computational Imaging*, vol. 7, pp. 849–863, 2021.

[10] U. S. Kamilov, C. A. Bouman, G. T. Buzzard, and B. Wohlberg, "Plug-and-play methods for integrating physical and learned models in computational imaging: Theory, algorithms, and applications," *IEEE Signal Processing Magazine*, vol. 40, no. 1, pp. 85–97, 2023.

[11] U. S. Kamilov, H. Mansour, and B. Wohlberg, "A plug-and-play priors approach for solving nonlinear imaging inverse problems," *IEEE Signal Processing Letters*, vol. 24, no. 12, pp. 1872–1876, 2017.

[12] Y. Sun, B. Wohlberg, and U. S. Kamilov, "An online plug-and-play algorithm for regularized image reconstruction," *IEEE Transactions on Computational Imaging*, vol. 5, no. 3, pp. 395–408, 2019.

[13] X. Xu, Y. Sun, J. Liu, B. Wohlberg, and U. S. Kamilov, "Provable convergence of plug-and-play priors with MMSE denoisers," *IEEE Signal Processing Letters*, vol. 27, pp. 1280–1284, 2020.

[14] Y. Romano, M. Elad, and P. Milanfar, "The little engine that could: Regularization by denoising (RED)," *SIAM Journal on Imaging Sciences*, vol. 10, no. 4, pp. 1804–1844, 2017.

[15] E. T. Reehorst and P. Schniter, "Regularization by denoising: Clarifications and new interpretations," *IEEE Transactions on Computational Imaging*, vol. 5, no. 1, pp. 52–67, 2019.

[16] R. Cohen, M. Elad, and P. Milanfar, "Regularization by denoising via fixed-point projection (RED-PRO)," *SIAM Journal on Imaging Sciences*, vol. 14, no. 3, pp. 1374–1406, 2021.

[17] H. W. Engl, "Discrepancy principles for Tikhonov regularization of ill-posed problems leading to optimal convergence rates," *Journal of optimization theory and applications*, vol. 52, pp. 209–215, 1987.

[18] L. Zanni, A. Benfenati, M. Bertero, and V. Ruggiero, "Numerical methods for parameter estimation in Poisson data inversion," *Journal of Mathematical Imaging and Vision*, vol. 52, no. 3, pp. 397–413, 2015.

[19] J. Immerkaer, "Fast noise variance estimation," *Computer vision and image understanding*, vol. 64, no. 2, pp. 300–302, 1996.

[20] S. Magnússon, P. C. Weeraddana, M. G. Rabbat, and C. Fischione, "On the convergence of alternating direction lagrangian methods for nonconvex structured optimization problems," *IEEE Transactions on Control of Network Systems*, vol. 3, no. 3, pp. 296–309, 2015.

[21] Y. Wang, W. Yin, and J. Zeng, "Global convergence of ADMM in nonconvex nonsmooth optimization," *Journal of Scientific Computing*, vol. 78, pp. 29–63, 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," in *2021 21st International Conference on Computational Science and Its Applications (ICCSA)*, 2021, pp. 39–46.

[23] G. Mataev, P. Milanfar, and M. Elad, "DeepRED: Deep image prior powered by RED," in *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2019, pp. 0–0.

[24] M. Bevilacqua, A. Roumy, C. Guillemot, and M. L. Alberi-Morel, "Low-complexity single-image super-resolution based on nonnegative neighbor embedding," *Proceedings of the 23rd British Machine Vision Conference (BMVC).*, 2012.

[25] "Kodak lossless true color image suite," https://r0k.us/graphics/kodak/.

[26] K. Zhang, W. Zuo, S. Gu, and L. Zhang, "Learning deep cnn denoiser prior for image restoration," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.