

On learning optimized reaction diffusion processes for effective image restoration

Yunjin Chen^{1,2}

Wei Yu¹

Thomas Pock^{1,3}

¹Graz University of Technology

²National University of Defense Technology

³Digital Safety & Security Department, AIT Austrian Institute of Technology GmbH

Abstract

For several decades, image restoration remains an active research topic in low-level computer vision and hence new approaches are constantly emerging. However, many recently proposed algorithms achieve state-of-the-art performance only at the expense of very high computation time, which clearly limits their practical relevance. In this work, we propose a simple but effective approach with both high computational efficiency and high restoration quality. We extend conventional nonlinear reaction diffusion models by several parametrized linear filters as well as several parametrized influence functions. We propose to train the parameters of the filters and the influence functions through a loss based approach. Experiments show that our trained nonlinear reaction diffusion models largely benefit from the training of the parameters and finally lead to the best reported performance on common test datasets for image restoration. Due to their structural simplicity, our trained models are highly efficient and are also well-suited for parallel computation on GPUs.

1. Introduction

Image restoration is the process of estimating uncorrupted images from noisy or blurred ones. It is one of the most fundamental operation in image processing, video processing, and low-level computer vision. There exists a huge amount of literature addressing the topic of image restoration problems, see for example [31] for a survey. Broadly speaking, most state-of-the-art techniques mainly concentrate on achieving utmost image restoration quality, with little consideration on the computational efficiency [44, 30, 19]. However, there are two notable exceptions, BM3D [11] and the recently proposed Cascade of Shrinkage Fields (CSF) [38] model, which simultaneously offer high efficiency and high image restoration quality.

It is well-known that BM3D is a highly engineered method, specialized for Gaussian noise. Moreover, it involves a block matching process, which is challenging for parallel computation on GPUs, alluding to the fact that it is not straightforward to accelerate BM3D algorithm on

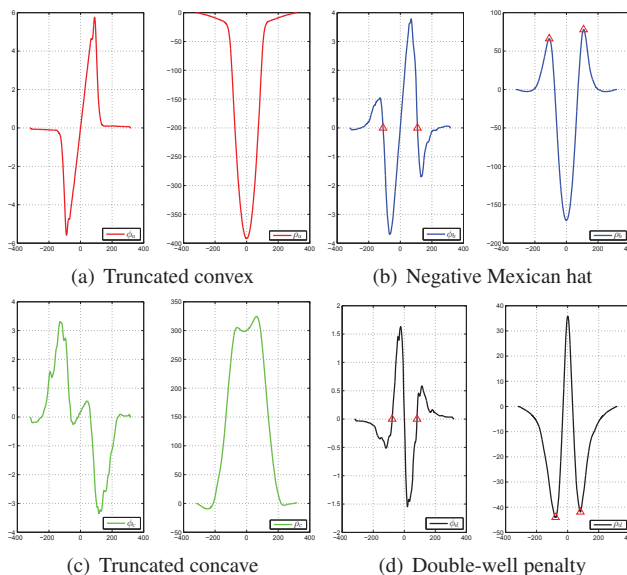


Figure 1. The figure shows four characteristic influence functions (left plot in each subfigure) together with their corresponding penalty functions (right plot in each subfigure), learned by our proposed method. A major finding in this paper is that our learned penalty functions significantly differ from the usual penalty functions adopted in partial differential equations and energy minimization methods. In contrast to their usual robust smoothing properties which is caused by a single minimum around zero, most of our learned functions have multiple minima different from zero and hence are able to enhance certain image structures. See Sec. 4.1 for more information.

parallel architectures. In contrast, the recently proposed CSF model offers high levels of parallelism, making it well suited for GPU implementation, thus owning high computational efficiency.

Among the approaches to tackle the problem of image restoration, nonlinear anisotropic diffusion [34, 41] defines a class of efficient approaches, as each diffusion step merely contains the convolution operation with a few linear filters. A nonlinear diffusion process usually corresponds to certain Partial Differential Equation (PDE) formulation. However, up to now, the image restoration quality of diffusion based approaches is still far away from the state-of-the-art, although with many improvements [21, 13, 35, 20].

We give a brief review of nonlinear diffusion based approaches and then introduce our proposed diffusion model. In the seminal work [34], Perona and Malik (P-M) demonstrated that nonlinear diffusion models yield very impressive results for image processing. This has given rise to many revised models with various formulations. A notable variant is the so-called biased anisotropic diffusion (also known as reaction diffusion) proposed by Nordström [32], which introduces a bias term (forcing term) to free the user from the difficulty of specifying an appropriate stopping time for the P-M diffusion process. This additional term reacts against the strict smoothing effect of the pure P-M diffusion, therefore resulting in a nontrivial steady-state.

Tsiotsios *et al.* [40] discussed the choice of some crucial parameters in the P-M model, such as the diffusivity function, the gradient threshold parameter and the stopping time of the iterative process. Some works consider modification to the diffusion term or the reaction term for the reaction diffusion model [15, 10, 32, 35, 1], *e.g.*, Acton *et al.* [1] and Plonkna *et al.* [35] exploited a more complicated reaction term to enhance oriented textures; [39, 3] proposed to replace the ordinary diffusion term with a flow equation based on mean curvature. Gilboa *et al.* [17] proposed a forward and backward diffusion process, which incorporates explicit inverse diffusion with negative diffusivity coefficient by carefully choosing the diffusivity function. The resultant diffusion processes can adaptively switch between forward and backward diffusion process. In a latter work [42], the theoretical foundations for discrete forward-and-backward diffusion filtering were investigated. Researchers also propose to exploit higher-order nonlinear diffusion filtering, which involves larger linear filters, *e.g.*, fourth-order diffusion models [21, 13, 20]. Meanwhile, theoretical properties about the stability and local feature enhancement of higher-order nonlinear diffusion filtering are established in [12].

It should be noted that all the above mentioned diffusion processes are handcrafted models. It is a generally difficult task to design a good-performing PDE for a specific image processing problem because good insights into this problem and a deep understanding of the behavior of the PDEs are usually required. Therefore, some researcher propose to learn PDEs from training data via an optimal control approach [29]. Unfortunately, at present [29] is the sole previous work we can find in this direction. The basic idea of our approach is the similar to [29], but we go much further and our proposed model is much more expressive.

1.1. Motivation and Contributions

In this paper we focus on nonlinear diffusion process due to its high efficiency and propose a trainable nonlinear diffusion model, which is parameterized by the linear filters and the influence functions. The trained diffusion model

contains many special influence functions (see Fig. 1 for an illustration), which greatly differ from usual influence functions employed in conventional diffusion models. It turns out that the trained diffusion processes can lead to effective image restoration with state-of-the-art performance, while preserve the property of high efficiency of diffusion based approaches. At present, we are not aware of any previous works that simultaneously optimize the linear filters and influence functions of a nonlinear diffusion process¹.

Our proposed nonlinear diffusion process has several remarkable benefits as follows:

- 1) It is conceptually simple as it is just a time-dynamic nonlinear reaction diffusion model with trained filters and influence functions;
- 2) It has broad applicability to a variety of image restoration problems. In principle, all existing diffusion based models can be revisited with appropriate training;
- 3) It yields excellent results for several tasks in image restoration, including Gaussian image denoising, and JPEG deblocking;
- 4) It is computationally very efficient and well suited for parallel computation on GPUs.

2. Proposed reaction diffusion process

We start with conventional nonlinear diffusion processes, then propose a training based reaction diffusion model for image restoration. Finally we show the relations between the proposed model and existing image restoration models.

2.1. Perona and Malik diffusion model

In the whole paper, we stick to the fully discrete setting, where images are represented as column vectors, *i.e.*, $u \in \mathbb{R}^N$. Therefore, the discrete version of the well-known Perona-Malik type nonlinear diffusion process [34] can be formulated as the following discrete PDE with an explicit finite difference scheme

$$\frac{u_{t+1} - u_t}{\Delta t} = - \sum_{i=\{x,y\}} \nabla_i^\top \Lambda(u_t) \nabla_i u_t \doteq - \sum_{i=\{x,y\}} \nabla_i^\top \phi(\nabla_i u_t), \quad (1)$$

where matrices ∇_x and $\nabla_y \in \mathbb{R}^{N \times N}$ are finite difference approximation of the gradient operators in x -direction and y -direction, respectively and Δt denotes the time step. $\Lambda(u_t) \in \mathbb{R}^{N \times N}$ is defined as a diagonal matrix

$$\Lambda(u_t) = \text{diag} \left(g \left(\sqrt{(\nabla_x u_t)_p^2 + (\nabla_y u_t)_p^2} \right) \right)_{p=1, \dots, N},$$

where function g is known as edge-stopping function [4] or diffusivity function [41], a typical g function given as

¹ Even though the linear filters and penalty functions in the image prior model [36, 8] can be trained simultaneously, the penalty function is optimized only in the sense that the weight α of certain fixed function (*e.g.*, $\alpha \cdot \log(1 + z^2)$) can be tuned. Our approach can exploit more generalized penalty functions (actually arbitrary functions), which is intractable in those previous models.

$g(z) = 1/(1+z^2)$. If ignoring the coupled relation between $\nabla_x u$ and $\nabla_y u$, the P-M model can be also written as the second formula on the right side in (1), where $\phi(\nabla_i u) = (\phi(\nabla_i u)_1, \dots, \phi(\nabla_i u)_N)^\top \in \mathbb{R}^N$ with function $\phi(z) = zg(z)$, known as influence function [4] or flux function [41]. In the upcoming subsection, we will stick to this decoupled formulation.

2.2. Proposed nonlinear diffusion model

Clearly, the matrix-vector product, $\nabla_x u$ can be interpreted as a 2D convolution of u with the linear filter $k_x = [-1, 1]$ (∇_y corresponds to the linear filter $k_y = [-1, 1]^\top$). Intuitively, in order to improve the capacity of the diffusion model, we can employ more filters of larger kernel size, in contrast to previous works that typically involve few filters with relatively small kernel size. We can additionally consider different influence functions for different filters, rather than a unique one. Moreover, the parameters of each iteration can vary across iterations. Taking the reaction term into account, our proposed nonlinear reaction diffusion model is formulated as

$$\frac{u_t - u_{t-1}}{\Delta t} = - \underbrace{\sum_{i=1}^{N_k} K_i^{t\top} \phi_i^t(K_i^t u_{t-1})}_{\text{diffusion term}} - \underbrace{\psi(u_{t-1}, f_n)}_{\text{reaction term}}, \quad (2)$$

where $K_i \in \mathbb{R}^{N \times N}$ is a highly sparse matrix, implemented as 2D convolution of the image u with the filter kernel k_i , i.e., $K_i u \Leftrightarrow k_i * u$, K_i is a set of linear filters and N_k is the number of filters. Function ϕ_i operates point-wise to the filter response $K_i u$. In practice, we set $\Delta t = 1$, as we can freely scale the formula on the right side. Note that in our proposed diffusion model, the influence functions are adjustable and can be different from each other.

The specific formulation for the reaction term $\psi(u)$ depends on applications. For classical image restoration problems, such as Gaussian denoising, image deblurring, image super resolution and image inpainting, we can set the reaction term to be the gradient of a data term, i.e. $\psi(u) = \nabla_u \mathcal{D}(u)$. For example, if $\mathcal{D}(u, f_n) = \frac{\lambda}{2} \|Au - f_n\|_2^2$, $\psi(u) = \lambda A^\top (Au - f_n)$, where f_n is the degraded image, A is the associated linear operator, and λ is related to the strength of the reaction term. In the case of Gaussian denoising, A is the identity matrix.

In our work, instead of making use of the well-chosen filters and influence functions, we train the nonlinear diffusion process for specific image restoration problem, including both the linear filters and the influence functions. As the diffusion process is an iterative approach, typically we run it for certain iterations. In order to make our proposed diffusion process more flexible, we train the parameters of the diffusion model for each single iteration. Finally, we arrive at a diffusion process which merely involves several iterations (referred to as stages).

2.3. Relations to existing image restoration models

Previous works [37, 32] show that in the nonlinear diffusion framework, there exist natural relations between reaction diffusion and regularization based energy functional. First of all, we can interpret (2) as one gradient descent step at u_{t-1} of a certain energy functional given by

$$E(u, f_n) = \sum_{i=1}^{N_k} \mathcal{R}_i(u) + \mathcal{D}(u, f_n), \quad (3)$$

where $\mathcal{R}_i(u) = \sum_{p=1}^N \rho_i^t((K_i^t u)_p)$ are the regularizers and the functions ρ_i^t are the so-called penalty functions. Note that $\rho'(z) = \phi(z)$. Since the parameters $\{K_i^t, \rho_i^t\}$ vary across the stages, (3) is a dynamic energy functional, which changes at each iteration.

In the case of fixed $\{K_i^t, \rho_i^t\}$ across the stages t , it is obvious that functional (3) is exactly the fields of experts image prior regularized variational model for image restoration [36, 9, 8]. In our work, we do not exactly solve this minimization problem anymore, but in contrast, we run the gradient descent step for several stages, and each gradient descent step is optimized by training.

In a very recent work [38], Schmidt *et al.* exploited an additive form of half-quadratic optimization to solve the same problem (3), which finally leads to a fast and effective image restoration model called cascade of shrinkage fields (CSF). The CSF model makes an assumption that the data term in (3) is quadratic and the operator A can be interpreted as a convolution operation, such that the corresponding subproblem has fast closed-form solution based on discrete Fourier transform (DFT). This restrains its applicability to many other problems such as image super resolution. However, our proposed diffusion model does not have this restriction on the data term. In principle, any smooth data term is appropriate. Moreover, as shown in the following sections, we can even handle the case of non-smooth data term.

There exist some previous works [2, 14], also trying to train an optimized gradient descent algorithm for the energy functional similar to (3). In their works, the Gaussian denoising problem is considered, and the trained gradient descent algorithm typically involves less than 10 iterations. However, their model is much more constrained, in the sense that, they exploited the same filters for each gradient descent step. More importantly, the influence function in their model is fixed to be a unique one. This clearly restricts the model performance, as demonstrated in Sec. 4.1.

There are also few preliminary works, e.g., [29] to go beyond traditional PDEs of the form (1), and propose to learn optimal PDEs for image restoration via optimal control. However, the investigated PDE model in [29] is too simple to generate a promising performance, as they only optimize the linear combination coefficients of a few predefined terms, which depend on selected derivative filters.

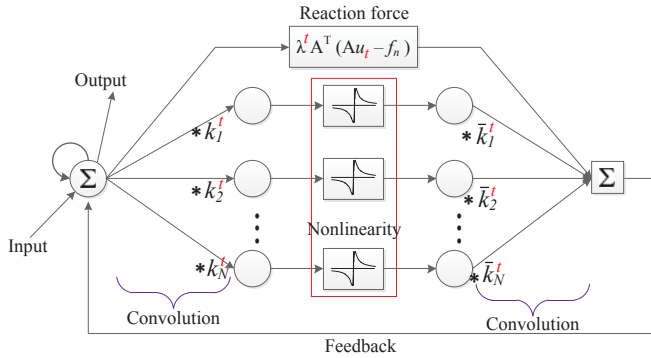


Figure 2. The architecture of our proposed diffusion model. Note that the additional convolution step with the rotated kernels \bar{k}_i (cf. Equ. 7) does not appear in conventional feed-forward CNs. Our model can be interpreted as a CN with a feedback step, which makes it different from conventional feed-forward networks. Due to the feedback step, it can be categorized into recurrent neural networks [18].

The proposed diffusion model also bears an interesting link to the convolutional networks (CNs) employed for image restoration problems [24]. One can see that each iteration (stage) of our proposed diffusion process involves the convolution operation with a set of linear filters, and thus it can be treated as a convolutional network. The architecture of our proposed network is shown in Figure 2, where one can see that it is not a pure feed-forward network any more, because it has a feedback step. Therefore, the structure of our CN model is different from conventional feed-forward networks. Due to this feedback step, it can be categorized into recurrent neural networks [18]. Moreover, the nonlinearity (e.g., influence functions in the context of nonlinear diffusion) in our proposed network are trainable. However, conventional CNs make use of fixed activation function, e.g., ReLU functions or sigmoid functions.

3. Learning

In this paper, we train our models for two representative image restoration problems: (1) denoising of images corrupted by Gaussian noise and (2) JPEG blocking artifacts reduction, which is formulated as a non-smooth problem. We use a loss minimization scheme to learn the model parameters $\Theta_t = \{\lambda^t, \phi_i^t, k_i^t\}$ for each stage t of the diffusion process, given S training samples $\{f_n^{(s)}, u_{gt}^{(s)}\}_{s=1}^S$, where $f_n^{(s)}$ is a noisy input and $u_{gt}^{(s)}$ is the corresponding ground truth clean image.

We firstly consider a greedy training strategy to train the diffusion processes stage-by-stage, i.e., at stage t , we minimize the cost function

$$\mathcal{L}(\Theta_t) = \sum_{s=1}^S \ell(u_t^{(s)}, u_{gt}^{(s)}), \quad (4)$$

where $u_t^{(s)}$ is the output of stage t of the diffusion process. We prefer the usual quadratic loss function to the negative PSNR used in [38], because the latter one imposes more

weights on those samples with relatively smaller cost, and thus leads to slightly inferior results in practice. The loss function is given as

$$\ell(u_t^{(s)}, u_{gt}^{(s)}) = \frac{1}{2} \|u_t^{(s)} - u_{gt}^{(s)}\|_2^2. \quad (5)$$

Parameterizing the influence functions ϕ_i^t : We parameterize the influence function via standard radial basis functions (RBFs), i.e., each function ϕ is represented as a weighted linear combination of a family of RBFs as follows

$$\phi_i^t(z) = \sum_{j=1}^M w_{ij}^t \varphi\left(\frac{|z - \mu_j|}{\gamma_j}\right), \quad (6)$$

where φ represents different RBFs. In this paper, we exploit RBFs with equidistant centers μ_j and unified scaling γ_j . We investigate two typical RBFs [23]: (1) Gaussian radial basis and (2) triangular-shaped radial basis.

In general, the Gaussian RBF can provide better approximation for generally smooth function than the triangular-shaped RBF with the same number of basis functions. However, the triangular-shaped RBF based function parameterization has the advantage of computational efficiency. More details can be found in the *supplemental material* [7]. In our work, we consider both function parameterization methods, but only present the results achieved based on the Gaussian RBF.

Training for denoising: According to the diffusion equation (2), for image denoising, the output of stage t is given as

$$u_t = u_{t-1} - \left(\sum_{i=1}^{N_k} \bar{k}_i^t * \phi_i^t(k_i^t * u_{t-1}) + \lambda^t (u_{t-1} - f_n) \right), \quad (7)$$

where we explicitly use a convolution kernel \bar{k}_i (obtained by rotating the kernel k_i 180 degrees) to replace the K_i^T for the sake of model simplicity².

Training for deblocking: Motivated by [5], we consider a new variational model for JPEG deblocking based on the FoE image prior model

$$\arg \min_u E(u) = \sum_{i=1}^{N_k} \rho_i(k_i * u) + \mathcal{I}_Q(Du), \quad (8)$$

where \mathcal{I}_Q is a indicator function over the set Q (quantization constraint set). In JPEG compression, information loss happens in the quantization step, where all possible values in the range $[d - 0.5, d + 0.5]$ (d is an integer) are quantized to a single number d . Given a compressed data, we only know d . Therefore, all possible values in the interval $[d - 0.5, d + 0.5]$ define a convex set Q which finally leads

²We use the symmetric boundary condition in our work. In this case, K_i^T can be interpreted as the convolution kernel \bar{k}_i only in the central region. Therefore, we actually slightly modify the original model.

to a box constraint. The sparse matrix $D \in \mathbb{R}^{N \times N}$ denotes the block DCT transform. We refer to [5] for more details.

We derive the diffusion process w.r.t the variational model (8) using the proximal gradient method [33], which reads as

$$u_t = D^\top \text{proj}_Q \left(D \left(u_{t-1} - \sum_{i=1}^{N_k} \bar{k}_i^t * \phi_i^t(k_i^t * u_{t-1}) \right) \right), \quad (9)$$

where $\text{proj}_Q(\cdot)$ denotes the orthogonal projection onto Q . More details can be found in the *supplemental material* [7].

Gradients: We minimize (4) with commonly used gradient based L-BFGS algorithm [28]. The gradient of the loss function at stage t w.r.t the model parameters Θ_t is computed using standard chain rule, given as

$$\frac{\partial \ell(u_t, u_{gt})}{\partial \Theta_t} = \frac{\partial u_t}{\partial \Theta_t} \cdot \frac{\partial \ell(u_t, u_{gt})}{\partial u_t}, \quad (10)$$

where $\frac{\partial \ell(u_t, u_{gt})}{\partial u_t} = u_t - u_{gt}$ is directly derived from (5), $\frac{\partial u_t}{\partial \Theta_t}$ is computed from (7) for the training of denoising task or (9) for the deblocking training, respectively. We do not present the derivatives for specific model parameters due to space limitation. All derivatives can be found in the *supplemental material* [7].

Joint training: In (4), each stage is trained greedily such that the output of each stage is optimized according to the loss function, regardless of the total stages T used in the diffusion process. A better strategy would be to jointly train all the stages simultaneously. The joint training task is formulated as

$$\mathcal{L}(\Theta_1, \dots, \Theta_T) = \sum_{s=1}^S \ell(u_T^{(s)}, u_{gt}^{(s)}), \quad (11)$$

where the loss function only depends on u_T (the output of the final stage T). The gradients of the loss function w.r.t Θ_t is given as

$$\frac{\partial \ell(u_T, u_{gt})}{\partial \Theta_t} = \frac{\partial u_T}{\partial \Theta_t} \cdot \frac{\partial u_{t+1}}{\partial u_t} \dots \frac{\partial \ell(u_T, u_{gt})}{\partial u_T},$$

which is the standard back-propagation technique widely used in the neural networks learning [26]. Compared with the greedy training, we additionally need to calculate $\frac{\partial u_{t+1}}{\partial u_t}$. All the derivations can be found in the *supplemental material* [7].

4. Experiments

We used the same 400 training images as [38], and cropped a 180×180 region from each image, resulting in a total of 400 training samples of size 180×180 , *i.e.*, roughly 13 million pixels.

We trained the proposed diffusion process with at most 8 stages to observe its saturation behavior after some stages.

We first greedily trained T stages of our model with specific model capacity, then conducted a joint training for the parameters of the whole T stages.

In our work, we mainly considered two trained reaction diffusion (TRD) models.

TRD $_{5 \times 5}^T$, Fully trained model with 24 filters of size 5×5 ,

TRD $_{7 \times 7}^T$, Fully trained model with 48 filters of size 7×7 ,

where TRD $_{m \times m}^T$ denotes a nonlinear diffusion process of stage T with filters of size $m \times m$. The filters number is $m^2 - 1$, if not specified.

Note that the calculation of the gradients of the loss function in (10) can be accomplished with convolution technique efficiently, even with a simple Matlab implementation. The training time varies greatly for different configurations. Important factors include (1) model capacity, (2) number of training samples, (3) number of iterations taken by the L-BFGS, and (4) number of Gaussian RBF kernels used for function approximation. We report below the most time consuming cases.

In training, computing the gradients $\frac{\partial \mathcal{L}}{\partial \Theta}$ with respect to the parameters of one stage for 400 images of size 180×180 takes about 35s (TRD $_{5 \times 5}$), 75s (TRD $_{7 \times 7}$) or 165s (TRD $_{9 \times 9}$) with Matlab implementation on a server with CPUs: Intel(R) Xeon E5-2680 @ 2.80GHz (eight parallel threads, 63 Gaussian RBF kernels for the influence function parameterization). We typically run 200 L-BFGS iterations for optimization. Therefore, the total training time, *e.g.*, for the TRD $_{7 \times 7}^5$ model is about $5 \times (200 \times 75) / 3600 = 20.8h$. *Code for learning and inference is available on the authors' homepage www.GPU4Vision.org.*

4.1. Image denoising experiments

We started with the training model of TRD $_{5 \times 5}^T$. We first considered the greedy scheme to train a diffusion process up to 8 stages (*i.e.*, $T \leq 8$), in order to observe the asymptotic behavior of the diffusion process. After the greedy training was completed, we conducted joint training for a diffusion model of certain stages (*e.g.*, $T = 5$), by simultaneously tuning the parameters in all stages.

We initialized the joint training with the parameters obtained from greedy training, as this is guaranteed not to decrease the training performance. In previous work [38], it is shown that joint training a model with filters of size 5×5 or larger hardly makes a difference relative to the result obtained by the greedy training. However, in our work we observed that joint training always improves the result of greedy training.

Note that for the models trained in the greedy manner, we can stop the inference at any stage, as its output of each stage is optimized. However, for the jointly trained models, we have to run T stages, as in this case only the output of the T^{th} stage is optimized.

Method	σ		St.	$\sigma = 15$	
	15	25		TRD _{5×5}	TRD _{7×7}
BM3D	31.08	28.56	2	31.14	31.30
LSSC	31.27	28.70	5	31.30	31.42
EPLL	31.19	28.68	8	31.34	31.43
opt-MRF	31.18	28.66		$\sigma = 25$	
RTF ₅	–	28.75		TRD _{5×5}	TRD _{7×7}
WNNM	31.37	28.83	2	28.58	28.77
CSF _{5×5} ⁵	31.14	28.60	5	28.78	28.92
CSF _{7×7} ⁵	31.24	28.72	8	28.83	28.95

Table 1. Average PSNR (dB) on 68 images from [36] for image denoising with $\sigma = 15, 25$.

We first trained our diffusion models for the Gaussian denoising problem with standard deviation $\sigma = 25$. The noisy training images were generated by adding synthetic Gaussian noise with $\sigma = 25$ to the clean images. Once we have trained a diffusion model, we evaluated its performance on a standard test dataset of 68 natural images.³

We present the final results of the joint training in Table 1, together with a selection of recent state-of-the-art denoising algorithms, namely BM3D [11], LSSC [30], EPLL [44], opt-MRF [8], RTF₅ model [25] and two very recent methods: the CSF model [38] and WNNM [19]. We downloaded these algorithms from the corresponding author’s homepage, and used them as is.

Concerning the performance of our TRD_{5×5}^T models, we find that joint training usually leads to an improvement of about 0.1dB in the cases of $T \geq 5$. From Table 1, one can see that (1) the performance of the TRD_{5×5}^T model saturates after stage 5, *i.e.*, in practice, 5 stages are typically enough; (2) our TRD_{5×5}⁵ model has achieved significant improvement (28.78 *vs.* 28.60), compared to a similar model CSF_{5×5}⁵, which has the same model capacity and (3) moreover, our TRD_{5×5}⁸ model is on par with so far the best-reported algorithm - WNNM.

When comparing with some closely related models such as the FoE prior based variational model [8], the FoE derived CSF model [38] and convolutional networks (CNs) [24], our trained models can provide significantly superior performance. Therefore, a natural question arises: what is the critical factor in the effectiveness of the trained diffusion models? There are actually two main aspects in our training model: (1) the linear filters and (2) the influence functions. In order to have a better understanding of the trained models, we went through a series of experiments to investigate the impact of these two aspects.

Analysis of the proposed diffusion process: Concentrating on the model capacity of 24 filters of size 5×5 , we considered the training of a diffusion process with 10 steps, *i.e.*, $T = 10$ for the Gaussian denoising of noise level $\sigma = 25$. We exploited two main classes of configurations: (A) the

³The test images are strictly separate from the training datasets.

parameters of every stage are the same and (B) every diffusion stage is different from each other. In both configurations, we consider two cases: (I) only train the linear filters with fixed influence function $\phi(z) = 2z/(1+z^2)$ and (II) simultaneously train the filters and influence functions.

Based on the same training dataset and test dataset, we obtained the following results: (A.I) every diffusion step is the same, and only the filters are optimized with fixed influence function. This is a similar configuration to previous works [2, 14]. The trained model achieves a test performance of 28.47dB. (A.II) with additional tuning of the influence functions, the resulting performance is boosted to 28.60dB. (B.I) every diffusion step can be different, but only train the linear filters with fixed influence function. The corresponding model obtains the result of 28.56dB, which is equivalent to the variational model [8] with the same model capacity. Finally (B.II) with additional optimization of the influence functions, the trained model leads to a significant improvement with the result of 28.86dB.

The analysis experiments demonstrate that without the training of the influence functions, there is no chance to achieve significant improvements over previous works, no matter how hard we tune the linear filters. Therefore, we believe that the additional freedom to tune the influence functions is the critical factor of our proposed training model. After having a closer look at the learned influence functions of the TRD_{5×5}⁵ model, these functions reinforce our argument.

Learned influence functions: The form of 120 learned penalty functions ρ^4 in the TRD_{5×5}⁵ model can be divided into four classes (see the corresponding subfigures in Figure 1):

- Truncated convex penalty functions with low values around zero to encourage smoothness.
- Negative Mexican hat functions, which have a local minimum at zero and two symmetric local maxima.
- Truncated concave functions with smaller values at the two tails.
- Double-well functions, which have a local maximum (not a minimum any more) at zero and two symmetric local minima.

At first glance, the learned penalty functions (except (a)) differ significantly from the usually adopted penalty functions used in PDE and energy minimization methods. However, it turns out that they have a clear meaning for image regularization.

Regarding the penalty function (b), there are two critical points (indicated by red triangles). When the magnitude of the filter response is relatively small (*i.e.*, less than the critical points), probably it is stimulated by the noise and

⁴The penalty function $\rho(z)$ is integrated from the influence function $\phi(z)$ according to the relation $\phi(z) = \rho'(z)$

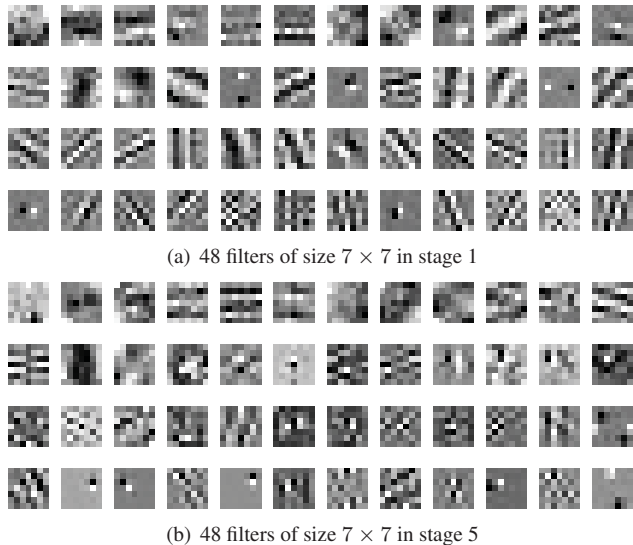


Figure 3. Trained filters (in the first and last stage) of the $\text{TRD}_{7 \times 7}^5$ model for the noise level $\sigma = 25$. We can find first, second and higher-order derivative filters, as well as rotated derivative filters along different directions. These filters are effective for image structure detection, such as image edge and texture.

therefore the penalty function encourages smoothing operation as it has a local minimum at zero. However, once the magnitude of the filter response is large enough (*i.e.*, across the critical points), the corresponding local patch probably contains a real image edge or certain structure. In this case, the penalty function encourages to increase the magnitude of the filter response, alluding to an image sharpening operation. Therefore, the diffusion process controlled by the influence function (b), can adaptively switch between image smoothing (forward diffusion) and sharpening (backward diffusion). We find that the learned influence function (b) is closely similar to an elaborately designed function in a previous work [17], which leads to an adaptive forward-and-backward diffusion process.

A similar penalty function to the learned function (c) with a concave shape is also observed in previous work on image prior learning [43]. This penalty function also encourages to sharpen the image edges. Concerning the learned penalty function (d), as it has local minima at two specific points, it prefers specific image structure, implying that it helps to form certain image structure. We also find that this penalty function is exactly the type of bimodal expert functions for texture synthesis employed in [22].

Now it is clear that the diffusion process involving the learned influence functions does not perform pure image smoothing any more for image processing. In contrast, it leads to a diffusion process for adaptive image smoothing and sharpening, distinguishing itself from previous commonly used image regularization techniques.

Influence of initialization: Our training model is also a deep model with many stages (layers). It is well-known that

Method	256^2	512^2	1024^2	2048^2	3072^2
BM3D [11]	1.1	4.0	17	76.4	176.0
CSF $_{7 \times 7}^5$ [38]	3.27	11.6	40.82	151.2	494.8
WNNM [19]	122.9	532.9	2094.6	–	–
$\text{TRD}_{5 \times 5}^5$	0.51	1.53	5.48	24.97	53.3
	0.43	0.78	2.25	8.01	21.6
	0.005	0.015	0.054	0.18	0.39
$\text{TRD}_{7 \times 7}^5$	1.21	3.72	14.0	62.2	135.9
	0.56	1.17	3.64	13.01	30.1
	0.01	0.032	0.116	0.40	0.87

Table 2. Run time comparison for image denoising (in seconds) with different implementations. (1) The run time results with gray background are evaluated with the single-threaded implementation on Intel(R) Xeon(R) CPU E5-2680 v2 @ 2.80GHz; (2) the blue colored run times are obtained with multi-threaded computation using Matlab *parfor* on the above CPUs; (3) the run time results colored in red are executed on a NVIDIA GeForce GTX 780Ti GPU. We do not count the memory transfer time between CPU/GPU for the GPU implementation (if counted, the run time will nearly double)

deep models are usually sensitive to initialization. However, our training model is not very sensitive to initialization. We have training experiments with fully random initializations for (1) greedy training. Using fully random initial parameters in range $[-0.5, 0.5]$, the trained models lead to a deviation within 0.01dB in the test phase and (2) joint training. Fully random initializations lead to models with inferior results, e.g., $\text{TRD}_{5 \times 5}^5$ (28.61 vs. 28.78). However, a plain initialization (all stages with DCT filters, influence function $\phi(z) = 2z/(1+z^2)$) works almost the same, e.g., $\text{TRD}_{5 \times 5}^5$, (28.75 vs. 28.78) and $\text{TRD}_{7 \times 7}^5$, (28.91 vs. 28.92).

Training for other configurations: In order to investigate the influence of the model capacity, we increase the filter size to 7×7 and 9×9 . We find that increasing the filter size from 5×5 to 7×7 brings a significant improvement of 0.14dB ($\text{TRD}_{7 \times 7}^5$ vs. $\text{TRD}_{5 \times 5}^5$) as show in Table 1. However, if we further increase the filter size to 9×9 , the resulting $\text{TRD}_{9 \times 9}^5$ leads to a performance of 28.96dB (a slight improvement of 0.05dB relative to the $\text{TRD}_{7 \times 7}^5$ model). In practice, we prefer the $\text{TRD}_{7 \times 7}^5$ model as it provides the best trade-off between performance and computation time. Fig. 3 shows the trained filters of the $\text{TRD}_{7 \times 7}^5$ model in the first and last stage.

We also trained diffusion models for the noise level of $\sigma = 15$, and the test performance is shown in Table 1. In experiments, we observed that joint training can always gain an improvement of about 0.1dB over the greedy training for the cases of $T \geq 5$.

From Table 1, one can see that for both noise levels, the resulting $\text{TRD}_{7 \times 7}^5$ model achieves the highest average PSNR. The $\text{TRD}_{7 \times 7}^5$ model outperforms the benchmark BM3D method by 0.35dB in average. This is a notable im-

q	JPEG decoder	TGV [5]	Dic-SR[6]	SADCT [16]	RTF[25]	TRD $_{7 \times 7}^4$
10	26.59	26.96	27.15	27.43	27.68	27.85
20	28.77	29.01	29.03	29.46	29.83	30.06
30	30.05	30.25	30.13	30.67	31.14	31.41

Table 3. JPEG deblocking results for natural images, reported with average PSNR values.

provement as few methods can surpass BM3D more than 0.3dB in average [27]. Moreover, the TRD $_{7 \times 7}^5$ model also surpasses the best-reported algorithm - WNNM method, which is very slow as shown in Table 2. In summary, our TRD $_{7 \times 7}^5$ model outperforms all the recent state-of-the-arts on the exploited test dataset, meanwhile it is the fastest method even with the CPU implementation.

Run time: The algorithm structure of our TRD model is closely similar to the CSF model, which is well-suited for parallel computation on GPUs. We implemented our trained models on GPU using CUDA programming to speed up the inference procedure, and finally it indeed lead to a significantly improved run time, see Table 2. We see that for the images of size up to $3K \times 3K$, the TRD $_{7 \times 7}^5$ model is still able to accomplish the denoising task in less than 1s.

We make a run time comparison to other denoising algorithms based on strictly enforced single-threaded CPU computation (*e.g.*, start Matlab with `-singleCompThread`) for a fair comparison, see Table 2. We only present the results of some selective algorithms, which either have the best denoising result or run time performance. We refer to [38] for a comprehensive run time comparison of various algorithms⁵.

We see that our TRD model is generally faster than the CSF model with the same model capacity. It is reasonable, because in each stage the CSF model involves additional DFT and inverse DTF operations, *i.e.*, our model only requires a portion of the computation of the CSF model. Even though the BM3D is a non-local model, it still possesses high computation efficiency. In contrast, another non-local model - WNNM achieves compelling denoising results but at the expense of huge computation time. Moreover, the WNNM algorithm is hardly applicable for high resolution images (*e.g.*, 10 mega-pixels) due to its huge memory requirements. Note that our model can be also easily implemented with multi-threaded CPU computation.

4.2. JPEG deblocking experiments

We also trained diffusion models for the JPEG deblocking problem. We followed the test procedure in [25] for performance evaluation. The test images were converted to gray-value, and scaled by a factor of 0.5, resulting images of

⁵ LSSC, EPLL, opt-MRF and RTF₅ methods are much slower than BM3D on the CPU, *cf.* [38].

size 240×160 . We distorted the images by JPEG blocking artifacts. We considered three compression quality settings $q = 10, 20$ and 30 for the JPEG encoder.

We trained three nonlinear diffusion TRD $_{7 \times 7}$ models for different compression parameter q . We found that for JPEG deblocking, 4 stages are already enough. Results of the trained models are shown in Table 3, compared with several representative deblocking approaches. We see that our trained TRD $_{7 \times 7}^4$ outperforms all the competing approaches in terms of PSNR. Furthermore, our model is extremely fast on GPU, *e.g.*, for a common image size of 1024×1024 , our model takes about 0.095s, while the strongest competitor (in terms of run time) - SADCT consumes about 56.5s with CPU computation⁶. See the *supplemental material* [7] for JPEG deblocking examples.

5. Conclusion and future work

We have proposed a trainable reaction diffusion model for effective image restoration. Its critical point lies in the training of the influence functions. We have trained our models for the problem of Gaussian denoising and JPEG deblocking. Based on standard test datasets, the trained models result in the best-reported results. We believe that the effectiveness of the trained diffusion models is attributed to the following desired properties of the models

- *Anisotropy.* In the trained filters, we can find rotated derivative filters in different directions, *cf.* Fig 3.
- *Higher order.* The learned filters contain first, second and higher-order derivative filters, *cf.* Fig 3.
- *Adaptive forward/backward diffusion through the learned nonlinear functions.* Nonlinear functions corresponding to explicit backward diffusion appear in the learned nonlinearity, *cf.* Fig 1.

Meanwhile, the trained models are very simple and well-suited for parallel computation on GPUs. As a consequence, the resulting algorithms are significantly faster than all competing algorithms and hence are also applicable to the restoration of high resolution images.

Future work: From an application point of view, we think that it will be interesting to consider learned, nonlinear reaction diffusion based models also for other image processing tasks such as image super resolution, blind image deconvolution, optical flow. Moreover, since learning the influence functions turned out to be crucial, we believe that learning optimal nonlinearities in CNs could lead to a similar performance increase. Finally, it will also be interesting to investigate the unconventional penalty functions learned by our approach in usual energy minimization approaches.

⁶ RTF is slower than SADCT, as it depends on the output of SADCT.

6. Acknowledgments

This work was supported by the Austrian Science Fund (FWF) under the China Scholarship Council (CSC) Scholarship Program and the START project BIVISION, No. Y729.

References

- [1] S. T. Acton, D. Prasad Mukherjee, J. P. Havlicek, and A. Conrad Bovik. Oriented texture completion by AM-FM reaction-diffusion. *IEEE TIP*, 10(6):885–896, 2001. 2
- [2] A. Barbu. Training an active random field for real-time image denoising. *IEEE TIP*, 18(11):2451–2462, 2009. 3, 6
- [3] C. A. Z. Barcelos, M. Boaventura, and E. C. Silva Jr. A well-balanced flow equation for noise removal and edge detection. *IEEE TIP*, 12(7):751–763, 2003. 2
- [4] M. Black, G. Sapiro, D. Marimont, and D. Heeger. Robust anisotropic diffusion and sharpening of scalar and vector images. In *ICIP*, pages 263–266. IEEE, 1997. 2, 3
- [5] K. Bredies and M. Holler. Artifact-free jpeg decompression with total generalized variation. In *VISAPP (I)*, pages 12–21, 2012. 4, 5, 8
- [6] H. Chang, M. K. Ng, and T. Zeng. Reducing artifact in JPEG decompression via a learned dictionary. *IEEE TSP*, 62(3):718–728, 2014. 8
- [7] Y. Chen. Notes on diffusion networks. *arXiv preprint arXiv:1503.05768*, 2015. 4, 5, 8
- [8] Y. Chen, T. Pock, R. Ranftl, and H. Bischof. Revisiting loss-specific training of filter-based mrfs for image restoration. In *GCPR*, pages 271–281, 2013. 2, 3, 6
- [9] Y. Chen, R. Ranftl, and T. Pock. Insights into analysis operator learning: From patch-based sparse models to higher order MRFs. *IEEE TIP*, 23(3):1060–1072, 2014. 3
- [10] G.-H. Cottet and L. Germain. Image processing through reaction combined with nonlinear diffusion. *Mathematics of Computation*, pages 659–673, 1993. 2
- [11] K. Dabov, A. Foi, V. Katkovnik, and K. O. Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *IEEE TIP*, 16(8):2080–2095, 2007. 1, 6, 7
- [12] S. Didas, J. Weickert, and B. Burgeth. Stability and local feature enhancement of higher order nonlinear diffusion filtering. In *DAGM*, pages 451–458. 2005. 2
- [13] S. Didas, J. Weickert, and B. Burgeth. Properties of higher order nonlinear diffusion filtering. *JMIV*, 35(3):208–226, 2009. 1, 2
- [14] J. Domke. Generic methods for optimization-based modeling. In *AISTATS*, pages 318–326, 2012. 3, 6
- [15] J. Esclarcín and L. Alvarez. Image quantization using reaction-diffusion equations. *SIAP*, 57(1):153–175, 1997. 2
- [16] A. Foi, V. Katkovnik, and K. Egiazarian. Pointwise shape-adaptive DCT for high-quality denoising and deblocking of grayscale and color images. *IEEE TIP*, 16(5):1395–1411, 2007. 8
- [17] G. Gilboa, N. Sochen, and Y. Y. Zeevi. Forward-and-backward diffusion processes for adaptive image enhancement and denoising. *IEEE TIP*, 11(7):689–703, 2002. 2, 7
- [18] A. Graves and J. Schmidhuber. Offline handwriting recognition with multidimensional recurrent neural networks. In *NIPS*, pages 545–552, 2009. 4
- [19] S. Gu, L. Zhang, W. Zuo, and X. Feng. Weighted nuclear norm minimization with application to image denoising. In *CVPR*, 2014. 1, 6, 7
- [20] P. Guidotti and K. Longo. Two enhanced fourth order diffusion models for image denoising. *JMIV*, 40(2):188–198, 2011. 1, 2
- [21] M. R. Hajiaboli. An anisotropic fourth-order diffusion filter for image noise removal. *IJCV*, 92(2):177–191, 2011. 1, 2
- [22] N. Heess, C. K. I. Williams, and G. E. Hinton. Learning generative texture models with extended Fields-of-Experts. In *BMVC*, pages 1–11, 2009. 7
- [23] Y. H. Hu and J.-N. Hwang. *Handbook of neural network signal processing*. CRC press, 2010. 4
- [24] V. Jain and S. Seung. Natural image denoising with convolutional networks. In *NIPS*, pages 769–776, 2009. 4, 6
- [25] J. Jancsary, S. Nowozin, and C. Rother. Loss-specific training of non-parametric image restoration models: A new state of the art. In *ECCV*, pages 112–125, 2012. 6, 8
- [26] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 5
- [27] A. Levin and B. Nadler. Natural image denoising: Optimality and inherent bounds. In *CVPR*, pages 2833–2840. IEEE, 2011. 8
- [28] D. C. Liu and J. Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45(1):503–528, 1989. 5
- [29] R. Liu, Z. Lin, W. Zhang, and Z. Su. Learning PDEs for image restoration via optimal control. In *ECCV*, pages 115–128. 2010. 2, 3
- [30] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman. Non-local sparse models for image restoration. In *ICCV*, pages 2272–2279, 2009. 1, 6
- [31] P. Milanfar. A tour of modern image filtering: new insights and methods, both practical and theoretical. *Signal Processing Magazine, IEEE*, 30(1):106–128, 2013. 1
- [32] K. Niklas Nordström. Biased anisotropic diffusion: a unified regularization and diffusion approach to edge detection. *Image and Vision Computing*, 8(4):318–327, 1990. 2, 3
- [33] P. Ochs, Y. Chen, T. Brox, and T. Pock. iPiano: Inertial Proximal Algorithm for Nonconvex Optimization. *SIAM Journal on Imaging Sciences*, 7(2):1388–1419, 2014. 5
- [34] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE TPAMI*, 12(7):629–639, 1990. 1, 2
- [35] G. Plonka and J. Ma. Nonlinear regularized reaction-diffusion filters for denoising of images with textures. *IEEE TIP*, 17(8):1283–1294, 2008. 1, 2
- [36] S. Roth and M. J. Black. Fields of Experts. *IJCV*, 82(2):205–229, 2009. 2, 3, 6
- [37] O. Scherzer and J. Weickert. Relations between regularization and diffusion filtering. *JMIV*, 12(1):43–63, 2000. 3
- [38] U. Schmidt and S. Roth. Shrinkage fields for effective image restoration. In *CVPR*, 2014. 1, 3, 4, 5, 6, 7, 8
- [39] V. Surya Prasath and D. Vorotnikov. Weighted and well-balanced anisotropic diffusion scheme for image denoising and restoration. *Nonlinear Analysis: Real World Applications*, 17:33–46, 2014. 2
- [40] C. Tsotsios and M. Petrou. On the choice of the parameters for anisotropic diffusion in image processing. *Pattern Recognition*, 46(5):1369–1381, 2013. 2
- [41] J. Weickert. *Anisotropic diffusion in image processing*. Teubner Stuttgart, 1998. 1, 2, 3
- [42] M. Welk, G. Gilboa, and J. Weickert. Theoretical foundations for discrete forward-and-backward diffusion filtering. In *SSVM*, pages 527–538. 2009. 2
- [43] S. C. Zhu and D. Mumford. Prior learning and Gibbs reaction-diffusion. *IEEE TPAMI*, 19(11):1236–1250, 1997. 7
- [44] D. Zoran and Y. Weiss. From learning models of natural image patches to whole image restoration. In *ICCV*, pages 479–486, 2011. 1, 6