

Patch-based learning of adaptive Total Variation parameter maps for blind image denoising

Claudio Fantasia*

Politecnico di Torino

DAUIN

Turin, Italy

claudio.fantasia@polito.it

Luca Calatroni*

MaLGa Center, DIBRIS,

Università di Genova,

MMS, Istituto Italiano di Tecnologia,

Genoa, Italy

luca.calatroni@unige.it

Xavier Descombes

Université Côte d'Azur

INRIA, CNRS, I3S

Sophia Antipolis, France

xavier.descombes@inria.fr

Rim Rekik

Université Grenoble-Alpes

INRIA, CNRS

Grenoble, France

rim.rekik-dit-nekhili@inria.fr

Abstract—We consider a patch-based learning approach defined in terms of neural networks to estimate spatially adaptive regularisation parameter maps for image denoising with weighted Total Variation (TV) and test it to situations when the noise distribution is unknown. As an example, we consider situations where noise could be either Gaussian or Poisson and perform preliminary model selection by a standard binary classification network. Then, we define a patch-based approach where at each image pixel an optimal weighting between TV regularisation and the corresponding data fidelity is learned in a supervised way using reference natural image patches upon optimisation of SSIM and in a sliding window fashion. Extensive numerical results are reported for both noise models, showing significant improvement w.r.t. results obtained by means of optimal scalar regularisation.

Index Terms—image denoising, weighted Total Variation, deep-learning for hyperparameter estimation.

I. INTRODUCTION

Hybrid approaches in the field of image restoration and image reconstruction represent nowadays a popular compromise between the good theoretical and interpretability properties of classical model-based approaches and the effectiveness of deep-learning data-driven procedures. Prominent examples in this respect are, for instance, algorithmic unrolling [1] and plug-and-play procedures [2] aiming at incorporating deeply-learned ingredients (i.e., regularisation functionals, their proximal operators) within standard regularisation schemes.

The task of image denoising offers an enlightening playground for the development of hybrid procedures, given the number of datasets and competing methods available nowadays. The task there consists in computing a noise-free version $\mathbf{x} \in \mathbb{R}^n$ from a noisy image $\mathbf{y} \in \mathbb{R}^n$ under some prior knowledge of the type of noising process described by a (possibly non-linear) function $\mathcal{T} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ so that:

$$\mathbf{y} = \mathcal{T}(\mathbf{x}), \quad (1)$$

* This author carried out this project in the joint Inria and CNRS MORPHEME team, UMR7271, Laboratoire I3S, Sophia-Antipolis, France.

The authors acknowledge the support received by the ANR-21-CE48-0008 PRC MICROBLIND project and of the ANR-22-CE48-0010 JCJC project. LC acknowledges the financial support of the European Research Council (ERC StG grant MALIN, 101117133).

Classical examples are, for instance, additive white Gaussian noise where $\mathcal{T}(\mathbf{x}) = \mathbf{x} + \mathbf{e}$ where $\mathbf{e} \sim \mathcal{N}(0, \sigma^2 \mathbf{Id})$ and signal-dependent Poisson noise where $\mathcal{T}(\mathbf{x}) = \mathcal{P}(\alpha \mathbf{x} + \boldsymbol{\eta})$, where $\alpha > 0$, $\mathbf{x} \in \mathbb{R}_{\geq 0}^n$ and $\boldsymbol{\eta} \in \mathbb{R}_{> 0}^n$ is used to have a strictly positive parameter.

Following a standard Bayesian approach, a solution of (1) can be found by maximising the posterior probability density function $\pi(\mathbf{x}|\mathbf{y})$, which via the Bayes formula corresponds to solve

$$\arg \max_{\mathbf{x}} \pi(\mathbf{x}|\mathbf{y}) = \frac{\pi(\mathbf{y}|\mathbf{x})\pi(\mathbf{x})}{\pi(\mathbf{y})}. \quad (2)$$

where $\pi(\mathbf{y})$ is a constant as $\mathbf{y}$ is known. Under the assumption of negative log-concave priors and likelihood terms, we can equivalently consider the problem of minimizing the composite functional being the sum of a data fidelity $\mathcal{D}(\cdot; \mathbf{y})$ and a regularisation term $\mathcal{R}$ in the form of

$$\arg \min_{\mathbf{x}} \mu \mathcal{D}(\mathbf{x}; \mathbf{y}) + \mathcal{R}(\mathbf{x}), \quad (3)$$

where the choice of the data term and of the regulariser depend on the particular noise-distribution/a-priori information, respectively, available on the data/solution and the scalar hyperparameter $\mu > 0$ balances the effect of the two.

Two prominent choices corresponding to the negative logarithm of the Gaussian and Poisson log-likelihoods are the Gaussian least-square fidelity [3]:

$$\mathcal{D}_G(\mathbf{x}; \mathbf{y}) = \sum_{i=1}^n d_G(x_i; y_i) = \sum_{i=1}^n \frac{1}{2} (x_i - y_i)^2, \quad (4)$$

and the Poisson Kullback-Leibler (KL) divergence [4] defined for $\mathbf{x} \geq 0$:

$$\begin{aligned} \mathcal{D}_P(\mathbf{x}; \mathbf{y}) &= \text{KL}_{\eta}(\mathbf{x}; \mathbf{y}) \\ &= \sum_{i=1}^n d_P(x_i; y_i) = \sum_{i=1}^n y_i \log \left(\frac{y_i}{x_i + \eta} \right) + x_i - y_i, \end{aligned} \quad (5)$$

where we defined the 1D components d_G and d_P in view of the following modelling. Note that we do not consider the variance in the Gaussian log-likelihood term as it is taken into account in parameter μ . The regularisation term $\mathcal{R}$ encodes *a priori* assumptions on the desired solution (e.g. sparsity,

smoothness). A celebrated choice is the following smoothed version of the Total Variation (TV) regularisation [3] which favours piece-wise constant smoothing and enjoys convexity and edge-preserving behaviour. It reads:

$$\text{TV}_\epsilon(\mathbf{x}) := \sum_{i=1}^n \|(\mathbf{D}\mathbf{x})_i\|_{2,\epsilon} = \sum_{i=1}^n \sqrt{(\mathbf{D}_h\mathbf{x})_i^2 + (\mathbf{D}_v\mathbf{x})_i^2 + \epsilon^2}, \quad (6)$$

where $\mathbf{D}\mathbf{x} = [\mathbf{D}_h\mathbf{x}; \mathbf{D}_v\mathbf{x}]^T$ denotes the image gradient and where the parameter $\epsilon > 0$ makes (6) differentiable. Optimal selection strategies for the parameter μ in the context of TV-regularised image reconstruction problems typically rely on grid-search approaches, generalised cross validation/L-curve [5] and statistical estimators exploiting the prior knowledge on noise intensity [6].

Traditionally, the parameter μ is a scalar value, so that regularisation is applied uniformly across the entire image domain. More recent approaches (see, e.g., [7]) consider weighted, that is spatially-adaptive, variants of TV-based models where a pixel-dependent weighting between regularisation and data term is used to account for differences between textured and geometric features, thus avoiding unbalanced effects corresponding to a *global* selection of μ , see Fig. 1, where the regularisation parameter maps is computed in a patch-wise fashion and optimal weighting values are computed on an exemplar image corrupted by Gaussian noise with $\sigma^2 = 0.01$.




Fig. 1: Space-variant Gaussian denoising for Gaussian noise with $\sigma^2 = 0.01$ with some close-ups in correspondence of some different geometric/textured regions.

The extension of the aforementioned parameter selection strategies for the computation of such pixel-dependent parameter map is often computationally unfeasible due to the dimensionality of the map (as big as the desired image), hence statistical [8] and highly-parametrised deep-learning based [9] approaches have been proposed for its computation.

A. Contribution

We consider a patch-based learning approach for estimating adaptive, i.e. spatially-varying regularisation maps for the problem of image denoising with unknown noise distribution. The regularisation map μ_{map} is estimated in a supervised learning fashion where, for each noisy patch $\mathbf{y}_p$, the map $\mathbf{y}_p \mapsto \mu_i$ providing an estimate of the regularisation parameter to-be-located in the pixel $i = 1, \dots, n$ placed at the centre of $\mathbf{y}_p$ is parametrised by a neural network f_θ which, at inference time, operates in a sliding-window fashion, taking as inputs patches of fixed size. An example of the proposed map estimation can be appreciated in Fig. 2.

Noise model selection is performed using a standard binary classification network g_ψ able to distinguish between Gaussian and Poisson noise. Several numerical results are reported showing visual and statistical improvements w.r.t. non-adaptive regularisation models.

II. PATCH-BASED LEARNING OF REGULARISATION MAPS

We consider in the following an adaptive approach that assigns to each pixel $i \in \{1, \dots, n\}$ of a noisy image $\mathbf{y}$ corrupted by Gaussian/Poisson noise a regularisation parameter μ_i computed by means of a suitable training of a deep neural network leveraging information on a small patch centered in the pixel itself. More precisely, given a patch $\mathbf{y}_{p_i}$ extracted by $\mathbf{y}$ around $i \in \{1, \dots, n\}$, we consider the following patch-wise version of problem (3) suited to deal with Gaussian/Poisson which reads:

$$\arg \min_{\mathbf{x}} \mu_i (\delta \mathcal{D}_G(\mathbf{x}; \mathbf{y}_{p_i}) + (1 - \delta) \mathcal{D}_P(\mathbf{x}; \mathbf{y}_{p_i})) + \text{TV}_\epsilon(\mathbf{x}), \quad (7)$$

where $\delta \in \{0, 1\}$ identifies (upon suitable training) whether Gaussian/Poisson noise is observed in the image, so that a tailored data term is used for reconstruction. The parameter μ_i is the one we would like to estimate to construct the desired regularisation map. We now parametrise the map $\mathbf{y}_{p_i} \mapsto \mu_i$ using a deep neural network f_θ . In the following, we describe the supervised-learning strategy which, given a dataset of clean/noisy natural image patches $\{(\hat{\mathbf{x}}_p^k, \mathbf{y}_p^k)\}_{k=1}^K$ learn such map. At inference time, note that this map will be applied in a sliding-window fashion, see Fig. 2, to an unobserved image $\mathbf{y}$ and used for solving:

$$\arg \min_{\mathbf{x}} \sum_{i=1}^n f_\theta(\mathbf{y}_{p_i}) (\delta d_G(x_i; y_i) + (1 - \delta) d_P(x_i; y_i)) + \text{TV}_\epsilon(\mathbf{x}), \quad (8)$$

where thus an adaptive weighting of the data-term against the TV regularisation is computed at each image pixel by considering patches around it.

A. Image patch dataset & network architecture

As a training dataset, we considered 300 images extracted from the popular Berkeley greyscale image dataset (BSD300), covering a diverse range of subjects (from faces, to natural landscapes etc). We preliminary split the dataset into

200/50/50 images for training/validation and testing. Overlapping patches of size 32×32 were extracted, with a sliding step set to 16 pixels, either to the right or downward, thus letting the patches to overlap by 25% or 50%. From each extracted patch, several noisy patches generated as different realisations of fixed Gaussian/Poisson noise distributions were generated. To create labels, i.e. reference values $\tilde{\mu}_i$ to be used for training, problem (7) was solved repeatedly (see Section II-B). To facilitate the learning process, the outlier values that did not fall in the range $(0, 1.5 \times \text{IQR}]$ were removed. The overall samples in the dataset for Poisson denoising are 144,870, for Gaussian denoising are 165,300. Two separate neural networks were trained for Gaussian and Poisson denoising. Both architectures have approximately 2.8×10^6 model parameters, as detailed in Table I. Its layer architecture is based on the one of the well-known AlexNet, with modifications aimed at reducing the total number of parameters, so as to mitigate overfitting (via dropout layers). As a training loss the mean squared error (MSE) between the predicted μ and the reference $\tilde{\mu}$ was used. Adam optimizer was used with an initial learning rate of 10^{-5} and an exponential decay rate of factor $\gamma = 0.8$, along with a weight decay of parameter 10^{-4} . Training and validation were performed over 30 epochs with a batch size of 256.

Layer Type	Output Shape	Kernel / Units
Input Layer	(1, 32, 32)	-
Conv2D + BatchNorm + ReLU	(64, 32, 32)	$5 \times 5, 64$
MaxPool2D	(64, 16, 16)	2×2
Conv2D + BatchNorm + ReLU	(128, 16, 16)	$5 \times 5, 128$
MaxPool2D	(128, 8, 8)	2×2
Conv2D + BatchNorm + ReLU	(256, 8, 8)	$3 \times 3, 256$
MaxPool2D	(256, 4, 4)	2×2
Conv2D + BatchNorm + ReLU	(512, 4, 4)	$3 \times 3, 512$
MaxPool2D	(512, 2, 2)	2×2
Flatten + BatchNorm1D	2048	-
Fully Connected + ReLU	512	512
Dropout (0.25)	512	-
Fully Connected + ReLU	128	128
Dropout (0.25)	128	-
Fully Connected	1	1

TABLE I: Details of the denoising architecture.

B. Solving the TV problem and computing labels

To create labels, i.e. reference values $\tilde{\mu}_i$ to be used for training, problem (7) was solved repeatedly for each training patch $\mathbf{y}_p$ being a noisy version of $\tilde{\mathbf{x}}_p$. More precisely, an optimal parameter $\tilde{\mu}$ and the corresponding solution $\tilde{\mathbf{x}}_{\tilde{\mu}}$ was computed by optimising $\text{SSIM}(\tilde{\mathbf{x}}, \mathbf{x}_{\mu})$ for different values of μ chosen within $[\mu_{\min}, \mu_{\max}]$.

To reduce the computational costs, a golden-section algorithm was employed. A good range for the parameters was experimentally determined to be $[\mu_{\min}, \mu_{\max}] = [0.01, 240]$.

In the case of Gaussian noise, problem (7) was solved using accelerated gradient descent (AGD) with fixed step-size. For Poisson denoising, a further step projecting the solution within the interval $[0, 1]$ at each iteration was employed in combination with an adaptive (non-monotone) backtracking strategy [10]. This guarantees faster convergence in comparison with fixed step-size approaches which can be quite inefficient in the case of Poisson noise due to the huge Lipschitz constant of the gradient of the KL data term [11], which leads to small step sizes and slow convergence.

C. Noise model selection: architecture

For choosing the appropriate data term (and, consequently, the reconstruction pipeline/algorithm) for computing labels, a binary noise classification task (Gaussian VS. Poisson) was performed. For that, patches of size 64×64 were used. In comparison to the denoising setting, larger patch size allows here for richer contextual information for distinguishing between noise types. For training the classification network, each clean patch of the dataset was injected with noise of Gaussian/Poisson distribution with varying intensity levels with a total number of 37,800 samples. Specifically, noise intensities were set to $\sigma^2 = \{0.01, 0.02, 0.03\}$ for Gaussian noise and $\alpha = \{45, 30, 10\}$ for Poisson noise. While structurally similar to the architecture detailed in Table I, the classification network differs in some key aspects: it applies higher channel counts in earlier layers (i.e. bigger patches), it lacks a MaxPool2D layer and it ends with a softmax output layer for binary classification. The increased number of channels thus results in a greater number of units in the dense layers so that the total number of parameters employed is approximately 1.7×10^7 . Training was performed by minimising the cross-entropy loss between predicted and ground-truth classes. As far as the optimization is concerned, the same setting as the one described for training the parameter estimation network was employed.

III. NUMERICAL RESULTS

The proposed approach is thus divided into two sequential steps: first the noise classifier network is applied to compute whether Gaussian ($\delta = 1$) or Poisson ($\delta = 0$) noise is observed in the image at hand. Then such parameter is plugged into (8), so that the trained network f_{θ} computing pixel-wise the parameter map μ is applied in a sliding-window fashion. A comprehensive overview of our proposed pipeline is shown in Fig. 2. In the following we report the results obtained in correspondence of a Gaussian denoising task with fixed variance $\sigma^2 = 0.01$ and a Poisson denoising task with fixed parameter $\alpha = 30$ after noise model selection is performed.

A. Noise model selection: performance

The accuracy on the classification task is measured using the following metric

$$\text{NN fitness} = \frac{\text{Number of correctly classified samples}}{\text{Number of samples}}. \quad (9)$$




Fig. 2: The input noisy image is taken as input by a noise classifier to identify the noise type. Based on the classification result, the corresponding denoising routine is used to estimate a pixel-wise regularisation map in a sliding-window fashion.

During training the accuracy obtained on the test set is NN fitness = 0.97, showing good performance. At inference time, the average accuracy is evaluated over 50 test images, each injected with different noise intensities. Taking as accuracy measure a similar definition to (9) but measuring the ratio between correctly classified and total number of patches, we obtained the results in Table II, covering a range from low noise to higher noise levels.

Noise Type	Avg accuracy (%)
Gaussian ($\sigma^2 = 0.01$)	0.67
Gaussian ($\sigma^2 = 0.015$)	0.90
Gaussian ($\sigma^2 = 0.02$)	0.99
Gaussian ($\sigma^2 = 0.03$)	1.0
Poisson ($\alpha = 60$)	0.98
Poisson ($\alpha = 45$)	0.99
Poisson ($\alpha = 30$)	0.99
Poisson ($\alpha = 15$)	0.98

TABLE II: Average accuracy for different types of noise.

B. Parameter map estimation: performance

The fitness of the network on the parameter estimation task is measured using the R^2 score:

$$R^2 := 1 - \frac{\sum_{k=1}^K (\tilde{\mu}_k - \mu_k)^2}{\sum_{k=1}^K (\tilde{\mu}_k - \mathbb{E}(\tilde{\mu}_k))^2}, \quad (10)$$

where K is the the number of test samples, $\tilde{\mu}_k$ are the labels while $\mu_k = f_\theta(\tilde{y}_p^k)$ is the output of the model while tested on the k -th patch .

The R^2 score for Gaussian denoising is 0.67, while for Poisson denoising it is 0.50. Such lower performance could be due to the fact that the Poisson case is more challenging to be dealt with. Although R^2 scores are far from values close to 1, our

method still achieves good quality reconstruction, as the key factor to consider is the ability to distinguish between regions with varying texture rather than absolute accuracy.

To validate the performance of the proposed approach, comparisons were performed on a set of test images between the space-variant method (8) exploiting locally adaptive regularisation, and a scalar approach where a single parameter $\mu_i \equiv \mu$ is computed—via SSIM maximisation (see Sec. 2B)—on the entire image. The two approaches are compared in terms of computational time, PSNR and SSIM.

C. Gaussian denoising $\sigma^2 = 0.01$

We report in the following some results obtained in correspondence of a noisy image y corrupted with Gaussian noise with zero mean and variance $\sigma^2 = 0.01$ and show improvements w.r.t. the scalar approach, see Fig. (3). Improvements upon SSIM and PSNR metrics are visible compared to the scalar approach, showing, in particular texture enhancement and stronger smoothing in homogeneous areas. We tested the approach over a dataset of 50 validation images, observing an average SSIM improvement of 0.017 and PSNR improvement of 0.35 w.r.t. the scalar approach. In terms of computational times, inference required on average ~ 187 seconds of computational time, in comparison of ~ 183 seconds required for computing the solution of the scalar approach with, alongside, optimal scalar parameter selection via golden section which requires direct comparison with the ground truth image and is unfeasible in case of the proposed pixel-dependent parameter map computation.

D. Poisson denoising $\alpha = 30$

We performed similar tests in correspondence with a Poisson denoising task for images corrupted with a noise level of $\alpha = 30$. Fig. 4 shows the results obtained in correspondence with an heterogeneous image containing both geometrical and texture areas. For this task, our method method resulted in an average SSIM improvement of 0.02 and an average PSNR improvement of 0.02 with respect to the scalar approach. In terms of computational times, inference time to perform the adaptive parameter map for our approach required ~ 187 seconds, in comparison with computation of the optimal parameter for the scalar approach for the scalar case which required ~ 200 seconds.

IV. CONCLUSIONS

We proposed a patch-wise deep-learning approach for estimating adaptive (i.e., pixel-dependent) regularisation parameter maps for (smoothed) TV image denoising in the presence of either Gaussian or Poisson noise which is detected about preliminary model selection. Spatial adaptivity combined with tailored noise-model selection allows the model to better handle heterogeneous image structures (geometry VS. texture), thus improving global reconstruction quality while mitigating (over/under)-smoothing effects observed in scalar parameter selection, in terms of both SSIM and PSNR, while maintaining a reasonable computational cost. Extensions of this work shall




Fig. 3: Gaussian denoising $\sigma^2 = 0.01$ on a test image.




Fig. 4: Poisson denoising $\alpha = 30$ on a test image.

address the actual inverse problem setting where a forward operator (such as a convolution or a MRI one as in [9]) is incorporated in the modelling.

REFERENCES

- [1] V. Monga, Y. Li, and Y. C. Eldar, “Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing,” *IEEE Signal Processing Magazine*, vol. 38, no. 2, pp. 18–44, 2021.
- [2] 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.
- [3] L. I. Rudin, S. Osher, and E. Fatemi, “Nonlinear total variation based noise removal algorithms,” *Physica D: Nonlinear Phenomena*, vol. 60, no. 1, pp. 259–268, 1992.
- [4] M. Bertero, P. Boccacci, G. Desiderà, and G. Vicidomini, “Image deblurring with poisson data: from cells to galaxies,” *Inverse Problems*, vol. 25, p. 123006, nov 2009.
- [5] P. C. Hansen, *Discrete Inverse Problems: Insight and Algorithms*. USA: Society for Industrial and Applied Mathematics, 2010.
- [6] Y.-W. Wen and R. H. Chan, “Parameter selection for total-variation-based image restoration using discrepancy principle,” *IEEE Transactions on Image Processing*, vol. 21, no. 4, pp. 1770–1781, 2012.
- [7] M. Pragliola, L. Calatroni, A. Lanza, and F. Sgallari, “On and beyond total variation regularization in imaging: The role of space variance,” *SIAM Review*, vol. 65, no. 3, pp. 601–685, 2023.
- [8] L. Calatroni, A. Lanza, M. Pragliola, and F. Sgallari, “A flexible space-variant anisotropic regularization for image restoration with automated parameter selection,” *SIAM Journal on Imaging Sciences*, vol. 12, no. 2, pp. 1001–1037, 2019.
- [9] A. Kofler, F. Altekürer, F. Antaro Ba, C. Kolbitsch, E. Papoutsellis, D. Schote, C. Sirotenko, F. F. Zimmermann, and K. Papafitsoros, “Learning regularization parameter-maps for variational image reconstruction using deep neural networks and algorithm unrolling,” *SIAM Journal on Imaging Sciences*, vol. 16, no. 4, pp. 2202–2246, 2023.
- [10] L. Calatroni and A. Chambolle, “Backtracking strategies for accelerated descent methods with smooth composite objectives,” *SIAM journal on optimization*, vol. 29, no. 3, pp. 1772–1798, 2019.
- [11] Z. T. Harmany, R. F. Marcia, and R. M. Willett, “This is spiral-tap: Sparse poisson intensity reconstruction algorithms—theory and practice,” *IEEE Transactions on Image Processing*, vol. 21, no. 3, pp. 1084–1096, 2011.