SYNTHESIS VERSUS ANALYSIS IN PATCH-BASED IMAGE PRIORS

Mário A. T. Figueiredo

Instituto de Telecomunicações and Instituto Superior Técnico, University of Lisbon, Portugal mario.figueiredo@tecnico.ulisboa.pt

ABSTRACT

In global models/priors (for example, using wavelet frames), there is a well known analysis vs synthesis dichotomy in the way signal/image priors are formulated. In patch-based image models/priors, this dichotomy is also present in the choice of how each patch is modeled. This paper shows that there is another analysis vs synthesis dichotomy, in terms of how the whole image is related to the patches, and that all existing patch-based formulations that provide a global image prior belong to the analysis category. We then propose a synthesis formulation, where the image is explicitly modeled as being synthesized by additively combining a collection of independent patches. We formally establish that these analysis and synthesis formulations are not equivalent in general and that both formulations are compatible with analysis and synthesis formulations at the patch level. Finally, we present an instance of the alternating direction method of multipliers (ADMM) that can be used to perform image denoising under the proposed synthesis formulation, showing its computational feasibility. Rather than showing the superiority of the synthesis or analysis formulations, the contributions of this paper is to establish the existence of both alternatives, thus closing the corresponding gap in the field of patch-based image processing.

Index Terms— Image restoration, image denoising, patch-based models, patch-based priors, analysis versus synthesis.

1. INTRODUCTION

The use of patches in image processing can be seen as an instance of the "divide and conquer" principle: since it is admittedly very difficult to formulate a global prior/model for images, patch-based approaches use priors/models for patches (rather than whole images), the combination of which yields the desired image prior/model. To keep the discussion and formulation at their essential and focus on the image modelling aspects, we will concentrate on image denoising, arguably the quintessential image processing problem. Nevertheless, much of what will be presented below can be easily extended (at least in principle) to more general inverse problems.

There are basically two approaches to patch-based image denoising. In earlier methods [3], [5], [6], patches are extracted from the noisy image, then processed/denoised independently (or maybe even collaboratively, as in BM3D [5]), and finally returned to their original locations. Since the patches overlap (to avoid blocking artifacts), there are several estimates of each pixel, which are combined by some form of averaging (unweighted or weighted) This approach is also used in the nonlocal Bayesian method [9], and in methods based on Gaussian mixtures [13]. For a comprehensive review of these and related methods, see [9]. Arguably, a conceptual flaw of these methods is that they obtain patch estimates without explicitly taking into account that these will subsequently be combined additively. Consequently, although some of these methods achieve state-of-the-art results, they do not explicitly provide a global image prior/model.

A more recent class of approaches builds a global image model based on a function computed from image patches, thus does not treat them as independent, and takes into account that they are overlapping patches of the same image; this approach was initiated with the *expected patch log-likelihood* (EPLL) [15], and adopted by most of the recent work [4], [12]. These methods do not have the conceptual flaw pointed out in the previous paragraph and provide a coherent global image model. In the synthesis vs analysis dichotomy, this class of patch-based models can be seen as an analysis method (as explained below in detail). This paper shows that there exists the synthesis counterpart of EPLL; in other words, that the synthesis vs analysis dichotomy is also present in the way the whole image and the patches are related.

The analysis vs synthesis dichotomy in global image models/priors (*e.g.*, based on wavelet frames, or total variation) has been first formalized in [7], and further studied in [11]; more recently, it has been ported to patch-wise models [4]. To the best of our knowledge, this dichotomy has not been pointed out before concerning the way in which patch-level models/priors are used to build a global image models; that is precisely the central contribution of this paper.

The remaining sections of this paper are organized as follows. After reviewing the classical analysis/synthesis dichotomy in Section 2, we shown in Section 3 that the classical patch-based methods follow an analysis formulation. Section 4 then introduces a synthesis patch-based formulation, and its relationship with the analysis counterpart is established in Section 5. In Section 6, we present an ADMM algorithm to efficiently perform image denoising under the proposed synthesis formulation. Finally, Section 7 concludes the paper by referring to future work directions.

2. ANALYSIS VS SYNTHESIS IN IMAGE DENOISING

Before addressing patch-based models, we briefly review the analysis and synthesis global formulations of image denoising [7], where the goal is to estimate an unknown image $\boldsymbol{x} \in \mathbb{R}^N$ (*N* is the total number of pixels in \boldsymbol{x} , which is a vectorized version of the corresponding $\sqrt{N} \times \sqrt{N}$ image) from a noisy version thereof

$$\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{w},\tag{1}$$

where \boldsymbol{w} is a sample of a white Gaussian noise field of zero mean and (known) variance σ^2 , that is, $p(\boldsymbol{y}|\boldsymbol{x}) = \mathcal{N}(\boldsymbol{y}; \boldsymbol{x}, \sigma^2 \boldsymbol{I})$.

The classical approach to estimate x from y is to adopt (or learn) a prior p_X for the unknown x and seek a maximizer of the posterior density (a *maximum a posteriori*–MAP–estimate)

$$\hat{\boldsymbol{x}}_{\text{MAP}} \in \operatorname*{argmax}_{\boldsymbol{x}} p(\boldsymbol{x}|\boldsymbol{y}) = \operatorname*{argmin}_{\boldsymbol{x}} \frac{1}{2\sigma^2} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2 - \log p_X(\boldsymbol{x}),$$
 (2)

where $\frac{1}{2\sigma^2} \| \boldsymbol{y} - \boldsymbol{x} \|_2^2 = -\log p(\boldsymbol{y}|\boldsymbol{x})$, up to an irrelevant constant. The analysis and synthesis formulations build priors for \boldsymbol{x} as follows.

Analysis: in this formulation, the prior p_X takes the form

$$-\log p_X(\boldsymbol{x}) = \varphi(\boldsymbol{F}\boldsymbol{x}) + b, \qquad (3)$$

where b is a constant and F a linear (analysis) operator. In this case, the MAP estimate takes the form

$$\hat{\boldsymbol{x}}_{\text{MAP-A}} \in \operatorname*{argmax}_{\boldsymbol{x}} p(\boldsymbol{x}|\boldsymbol{y}) = \operatorname*{argmin}_{\boldsymbol{x}} \frac{1}{2\sigma^2} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2 + \varphi(\boldsymbol{F}\boldsymbol{x}),$$
(4)

where MAP-A stands for MAP-analysis.

Synthesis: here, the starting point is to assume that \boldsymbol{x} is linearly synthesized/represented according to $\boldsymbol{x} = \boldsymbol{G}\boldsymbol{\theta}$, where $\boldsymbol{\theta}$ is a vector of coefficients, and the prior is formulated on $\boldsymbol{\theta}$ rather than directly on \boldsymbol{x} ; once an estimate $\hat{\boldsymbol{\theta}}$ is obtained, the corresponding estimate of \boldsymbol{x} is simply $\hat{\boldsymbol{x}} = \boldsymbol{G}\hat{\boldsymbol{\theta}}$. In summary,

$$\hat{\boldsymbol{x}}_{\text{MAP-S}} = \boldsymbol{G}\hat{\boldsymbol{\theta}}, \text{ where } \hat{\boldsymbol{\theta}} \in \operatorname*{argmin}_{\boldsymbol{\theta}} \frac{1}{2\sigma^2} \|\boldsymbol{y} - \boldsymbol{G}\boldsymbol{\theta}\|_2^2 + \zeta(\boldsymbol{\theta}),$$
(5)

and
$$\zeta(\boldsymbol{\theta}) = -\log_{\Theta}(\boldsymbol{\theta})$$
 is a negative log-prior on $\boldsymbol{\theta}$.

One of the main distinguishing features of analysis and synthesis formulations is that in the former the object of the estimation procedure is the image itself, whereas in the latter, one estimates a representation from which the image estimate is synthesized.

3. PATCH-ANALYSIS FORMULATION

3.1. Expected Patch Log-likelihood

A central tool in most patch-based approaches is a collection of operators $\{P_m, m = 1, ..., M\}$ that extracts M patches from a given image of size $\sqrt{N} \times \sqrt{N}$; each P_m can be seen as a binary matrix with size $n \times N$, where n is the total number of pixels in each patch (assumed square, of size $\sqrt{n} \times \sqrt{n}$). The standard way of formulating a patch-based prior is by writing

$$p_X(\boldsymbol{x}) = \frac{1}{Z} \prod_{m=1}^M f(\boldsymbol{P}_m \boldsymbol{x}), \qquad (6)$$

where $f : \mathbb{R}^n \to \mathbb{R}_+$ is a function expressing the patch-wise prior distribution and Z is a normalizing constant. Function f may itself be a *probability density function* (pdf), *e.g.*, a GMM, in which case this prior is an instance of a so-called *product of experts* (PoE [8]). However, f does not need to be a pdf, as long as it takes non-negative values; in fact, (6) can also be seen as a *factor graph* model, where each factor corresponds to a patch and the all factors have the same function f [14]. Moreover, this prior is also equivalent to the formulation known as EPLL (*expected patch log-likleihood* [15]), although EPLL was not originally interpreted as a prior.

Given a noisy image y, a MAP estimate of x is given by

$$\hat{\boldsymbol{x}}_{\text{MAP-A}} \in \arg\min_{\boldsymbol{x}} \frac{1}{2\sigma^2} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2 - \sum_{m=1}^M \log f(\boldsymbol{P}_m \boldsymbol{x}).$$
(7)

3.2. Half Quadratic Splitting and ADMM

To tackle the large-scale optimization problem in (7), the so-called *half quadratic splitting* strategy replaces it with

$$\hat{\boldsymbol{x}}_{\text{MAP-A}} \in \arg\min_{\boldsymbol{x}} \min_{\boldsymbol{v}} \frac{1}{\sigma^2} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2 + \beta \sum_{m=1}^M \|\boldsymbol{v}_m - \boldsymbol{P}_m \boldsymbol{x}\|_2^2$$
$$-2 \sum_{m=1}^M \log f(\boldsymbol{v}_m), \qquad (8)$$

where $v = (v_m, m = 1, ..., M)$, which obviously becomes equivalent to (7) as $\beta \to \infty$ [15]. The optimization problem in (8) is tackled by alternating between minimizing with respect to v and x, while slowly increasing β . Other strategies for setting β have also been proposed [15].

An obvious alternative (as recently mentioned in [10]) is to reformulate (7) as a constrained problem

$$\hat{x}_{\text{MAP-A}} = \arg \min_{\boldsymbol{x}} \frac{1}{2\sigma^2} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2 - \sum_{m=1}^M \log f(\boldsymbol{v}_m) \quad (9)$$
subject to $\boldsymbol{v}_m = \boldsymbol{P}_m \boldsymbol{x}$, for $m = 1, ..., M$

and tackle it with ADMM (alternating direction method of multipliers) [2]. Of course, convergence of ADMM for this problem can only be guaranteed if the negative log factors $-\log f$ are convex; this is not the case if f is a GMM, but it is true if $-\log f$ is an ℓ_r norm or the r-th power thereof, e.g., $-\log f(\boldsymbol{v}_m) = \|\boldsymbol{v}_m\|_r^r$, with $r \ge 1$.

Examining (8) (with $\beta \to \infty$) or (9) reveals that this formulation seeks a consensus among the patches, in the sense that the several replicates of each pixels that exist in different patches are forced to agree on a common value for that pixel. In other words, the clean patches are not modelled as additively generating a clean image, and are merely used to write a joint prior p_X that factorizes across overlapping patches.

3.3. Identification as Analysis Formulation

The estimation criterion in (7) clearly falls in the analysis-type category [7], since it considers the image itself as the object to be estimated and as the argument of the prior. However, as a generative model for clean images, its meaning is not very clear, since it is not a trivial task to obtain samples from this distribution.

To obtain a more compact notation, let $\boldsymbol{P} : \mathbb{R}^N \to \mathbb{R}^{Mn}$ be the operator (an $Mn \times N$ matrix) that extracts the set of M patches, *i.e.*,

$$\boldsymbol{P}\boldsymbol{x} = \begin{bmatrix} \boldsymbol{P}_1 \boldsymbol{x} \\ \vdots \\ \boldsymbol{P}_M \boldsymbol{x} \end{bmatrix} \in \mathbb{R}^{Mn}.$$
 (10)

The prior p_X may then be written as

$$p_X(\boldsymbol{x}) \propto p_V(\boldsymbol{P}\boldsymbol{x}),$$
 (11)

where p_V denotes a density defined in \mathbb{R}^{Mn} according to

$$p_V(\boldsymbol{v}) = \prod_{m=1}^M f(\boldsymbol{v}_m). \tag{12}$$

With this notation, the MAP denoising problem (7) can be written as

$$\hat{\boldsymbol{x}}_{\text{MAP-A}} \in \arg\min_{\boldsymbol{x}} \frac{1}{2\sigma^2} \|\boldsymbol{x} - \boldsymbol{y}\|_2^2 - \log p_V(\boldsymbol{P}\boldsymbol{x}), \quad (13)$$

which clearly reveals the analysis nature of this formulation (see (4)).

3.4. Patch level Models

We stress that the analysis/synthesis dichotomy addressed in this paper concerns the way in which an image relates to its patches, not the way the patches themselves are modelled. In fact, the patch-analysis formulation just reviewed is compatible with a synthesis patch model, *e.g.*, one that models each patch $P_m x$ as linear combination of elements of some dictionary D, with coefficients γ_m equipped with some prior p_{Γ} (for example, a sparsity-inducing prior, as in [12]). Using the half quadratic splitting approach, the formulation becomes

$$\hat{\boldsymbol{x}}_{\text{MAP-A}} \in \arg\min_{\boldsymbol{x}} \min_{\boldsymbol{\gamma}} \quad \frac{1}{\sigma^2} \|\boldsymbol{y} - \boldsymbol{x}\|_2^2 + \beta \sum_{m=1}^M \|\boldsymbol{D}\boldsymbol{\gamma}_m - \boldsymbol{P}_m \boldsymbol{x}\|_2^2$$
$$-2 \sum_{m=1}^M \log p_{\Gamma}(\boldsymbol{\gamma}_m). \tag{14}$$

where $\gamma = (\gamma_m, m = 1, ..., M)$. Naturally, the patch-analysis formulation is also compatible with an analysis patch model [4], by using a patch prior of the form $f = \phi \circ A$, *i.e.*, $f(v) = \phi(Av)$, for some function $\phi : \mathbb{R}^s \to \mathbb{R}_+$ and matrix $A \in \mathbb{R}^{s \times n}$.

To identify the analysis/synthesis dichotomy in terms of how the patches are related to the underlying image, not in terms of how the patches are modelled, we refer to the formulation reviewed in this section as *patch-analysis* and to its synthesis counterpart (to be introduced in the next section) as *patch-synthesis*.

4. PATCH-SYNTHESIS FORMULATION

We now present the patch-synthesis formulation, which can be summarized as follows: the clean image is generated by additively combining a collection of patches; the patches themselves follow some probabilistic model, but are a priori mutually independent.

Consider a collection of patches $\{z_m \in \mathbb{R}^n, m = 1, ..., M\}$ and let image x be synthesized from these patches by combining them additively according to

$$\boldsymbol{x} = \sum_{m=1}^{M} \boldsymbol{Q}_m \boldsymbol{z}_m, \qquad (15)$$

where matrices $\boldsymbol{Q}_m \in \mathbb{R}^{N \times n}$ are such that they average the values in the several patches that contribute to a given pixel of \boldsymbol{x} .

A simple 1D example will help clarify the structure of the Q_m matrices. Consider that $\boldsymbol{x} = [x_1, x_2, x_3, x_4]^T$ is produced by combining all the consecutive 2-element patches (with periodic boundary conditions), which correspond to the subsets of components $\{1, 2\}$, $\{2, 3\}$, $\{3, 4\}$, and $\{4, 1\}$. That is, M = 4 and

$$\boldsymbol{Q}_{1} = \begin{bmatrix} 1/2 & 0\\ 0 & 1/2\\ 0 & 0\\ 0 & 0 \end{bmatrix}, \quad \boldsymbol{Q}_{2} = \begin{bmatrix} 0 & 0\\ 1/2 & 0\\ 0 & 1/2\\ 0 & 0 \end{bmatrix}, \quad (16)$$

$$\boldsymbol{Q}_{3} = \begin{bmatrix} 0 & 0\\ 0 & 0\\ 1/2 & 0\\ 0 & 1/2 \end{bmatrix}, \quad \boldsymbol{Q}_{4} = \begin{bmatrix} 0 & 1/2\\ 0 & 0\\ 0 & 0\\ 1/2 & 0 \end{bmatrix}.$$
(17)

Stacking all the patches in vector $\boldsymbol{z} \in \mathbb{R}^{Mn}$ and considering a matrix $\boldsymbol{Q} = [\boldsymbol{Q}_1 \cdots \boldsymbol{Q}_M] \in \mathbb{R}^{N \times (Mn)}$, the synthesis expression in (15) can be written compactly as

$$\boldsymbol{x} = \boldsymbol{Q}\boldsymbol{z}.\tag{18}$$

With the patches modelled as independent and identically distributed samples of some patch-wise pdf g, their joint log-prior is

$$\log p_Z(\boldsymbol{z}) = \sum_{m=1}^M \log g(\boldsymbol{z}_m).$$
(19)

Notice that obtaining samples from (19) is as simple as obtaining samples from g itself. Consequently, generating image samples under this synthesis model simply corresponds to generating samples from (19) and then multiplying them by Q. This is in contrast with the patch-analysis prior (6), where, even if f is a valid pdf, it is not trivial to obtain samples from p_X .

The resulting MAP denoising criterion can now be written as

$$\hat{\boldsymbol{x}}_{\text{MAP-S}} = \boldsymbol{Q}\hat{\boldsymbol{z}}, \text{ where } \hat{\boldsymbol{z}} \in \arg\min_{\boldsymbol{z}} \frac{1}{2\sigma^2} \|\boldsymbol{Q}\boldsymbol{z} - \boldsymbol{y}\|_2^2 - \log p_Z(\boldsymbol{z}).$$
(20)

As the patch-analysis model, the patch-synthesis formulation that we have just presented is compatible with both analysis and synthesis priors for the patches, and of course with any valid pdf for vectors in \mathbb{R}^n . An analysis formulation simply amounts to choosing a patch prior of the form $g = \psi \circ B$, *i.e.*, $g(\boldsymbol{z}_m) = \psi(\boldsymbol{B}\boldsymbol{z}_m)$, for some function $\psi : \mathbb{R}^s \to \mathbb{R}_+$ and matrix $\boldsymbol{B} \in \mathbb{R}^{s \times n}$.

In a synthesis formulation, each patch z_m is synthesized using some dictionary D as $z_m = D\gamma_m$, and the γ_m follow some prior p_{Γ} ; stacking all the patch coefficients in vector γ , we can write $z = \Delta \gamma$ (where Δ is a block-diagonal matrix with M replicas of D), thus the MAP denoising criterion becomes

$$\hat{\boldsymbol{x}}_{\text{MAP-S}} = \boldsymbol{Q} \boldsymbol{\Delta} \hat{\boldsymbol{\gamma}}, \qquad (21)$$

where

$$\hat{\boldsymbol{\gamma}} \in \arg\min_{\boldsymbol{\gamma}} \frac{1}{2\sigma^2} \|\boldsymbol{Q}\boldsymbol{\Delta}\boldsymbol{\gamma} - \boldsymbol{y}\|_2^2 - \log q_{\Gamma}(\boldsymbol{\gamma}).$$
 (22)

5. RELATIONSHIP BETWEEN THE ANALYSIS AND SYNTHESIS FORMULATIONS

Leaving aside for now the choice of the patch priors, let us focus on the relationship between formulations (13) and (20). The key observation underlying the relationship between these two formulations is that (assuming the patch structure in both formulations is the same)

$$\boldsymbol{Q} \boldsymbol{P} = \boldsymbol{I},\tag{23}$$

where I denotes the identity matrix, but in general

$$\boldsymbol{Q} \boldsymbol{P} \neq \boldsymbol{I}.$$
 (24)

In other words, Q is a left pseudo-inverse of P. To prove (23), simply notice that if a collection of patches is extracted from some image and then these patches are used to synthesize an image by averaging the overlapping pixels, an identical image is obtained. Of course, the converse is not true, in general: if an image is synthesized by averaging the overlapping pixels of a collection of patches, and then patches are extracted from the synthesized image, there is no guarantee that these patches are equal to the original ones, thus proving (24). In the trivial and uninteresting cases where the patches are singletons, or non-overlapping, we would have Q P = I.

As shown in [7], given the analysis formulation (13), an equivalent synthesis formulation is

$$\hat{\boldsymbol{x}}_{\text{MAP-A}} = \boldsymbol{Q}\hat{\boldsymbol{z}}, \text{ where } \hat{\boldsymbol{z}} \in \arg\min_{\boldsymbol{z}} \frac{1}{2\sigma^2} \|\boldsymbol{Q}\boldsymbol{z} - \boldsymbol{y}\|_2^2 - \log p_V(\boldsymbol{z}),$$

subject to $\boldsymbol{z} \in \text{range}(\boldsymbol{P}),$
(25)

where the constraint $z \in \text{range}(P)$ enforces z to be in the subspace spanned by the columns of P. Notice that this constraint corresponds to having a collection of patches extracted from some image, *i.e.*, it forces the patches to agree on the value of each shared pixel. Since (20) does not enforce this constraint, it is not, in general, equivalent to the patch-analysis formulation (13).

6. ADMM FOR PATCH-SYNTHESIS DENOISING

In this section, we derive an instance of ADMM to deal with (20). The first step is to rewrite it as a constrained problem,

$$\min_{\boldsymbol{z},\boldsymbol{u}} \frac{1}{2\sigma^2} \|\boldsymbol{Q}\boldsymbol{z} - \boldsymbol{y}\|_2^2 + \xi(\boldsymbol{u}),$$

subject to $\boldsymbol{z} = \boldsymbol{u},$ (26)

where

$$\xi(\boldsymbol{u}) = -\log p_Z(\boldsymbol{u}) = \sum_{m=1}^{M} \underbrace{(-\log g(\boldsymbol{u}_m))}_{\xi_m(\boldsymbol{u}_m)}$$
(27)

is the negative log-prior, or regularizer. ADMM for this problem takes the form

$$\begin{aligned} \boldsymbol{z}^{(t+1)} &= \arg \min_{\boldsymbol{z}} \ \frac{1}{2\sigma^2} \| \boldsymbol{Q} \boldsymbol{z} - \boldsymbol{y} \|_2^2 + \frac{\rho}{2} \| \boldsymbol{z} - \boldsymbol{u}^{(t)} - \boldsymbol{d}^{(t)} \|_2^2 \\ \boldsymbol{u}^{(t+1)} &= \arg \min_{\boldsymbol{u}} \ \xi(\boldsymbol{u}) + \frac{\rho}{2} \| \boldsymbol{z}^{(t+1)} - \boldsymbol{u} - \boldsymbol{d}^{(t)} \|_2^2 \\ \boldsymbol{d}^{(t+1)} &= \boldsymbol{d}^{(t)} + \boldsymbol{z}^{(t+1)} - \boldsymbol{u}^{(t+1)} \end{aligned}$$

The update equation for $\boldsymbol{u}^{(t+1)}$ is a denoising step. Due to the separability of the squared ℓ_2 norm and of ξ (see (27)), this can be separately solved with respect to each patch: for m = 1, ..., M,

$$\boldsymbol{u}_{m}^{(t+1)} = \operatorname*{argmin}_{\boldsymbol{u}_{m}} \ \xi_{m}(\boldsymbol{u}_{m}) + \frac{\rho}{2} \|\boldsymbol{z}_{m}^{(t+1)} - \boldsymbol{u}_{m} - \boldsymbol{d}_{m}^{(t)}\|_{2}^{2}, \quad (28)$$

which is simply the proximity operator of $(1/\rho)\xi_m$ [1], computed at $d_m^{(t)} - z_m^{(t+1)}$. In a Bayesian viewpoint, (28) corresponds to obtaining the MAP estimate of u_m from observations $d_m^{(t)} - z_m^{(t+1)}$, assuming additive white Gaussian noise of variance $1/\rho$ and a negative logprior ξ_m . If the ξ_m are convex, this MAP estimate is unique, due to the strict convexity of the quadratic term in (28).

Computing $z^{(t+1)}$ corresponds to solving an unconstrained quadratic problem, the solution being

$$\boldsymbol{z}^{(t+1)} = \left(\boldsymbol{Q}^{T}\boldsymbol{Q} + \sigma^{2}\rho\boldsymbol{I}\right)^{-1} \left(\boldsymbol{Q}^{T}\boldsymbol{y} + \sigma^{2}\rho\left(\boldsymbol{u}^{(t)} + \boldsymbol{d}^{(t)}\right)\right).$$
(29)

The bottleneck in this update equation seems to be the matrix inversion, since $Q^T Q$ is a huge $(Mn) \times (Mn)$ matrix. However, this inversion can be solved very efficiently by resorting to the Sherman-Morrison-Woodbury matrix inversion formula. In fact,

$$\left(\boldsymbol{Q}^{T}\boldsymbol{Q}+\sigma^{2}\rho\,\boldsymbol{I}\right)^{-1}=\frac{1}{\sigma^{2}\rho}\left(\boldsymbol{I}-\boldsymbol{Q}^{T}(\sigma^{2}\rho\,\boldsymbol{I}+\boldsymbol{Q}\,\boldsymbol{Q}^{T})^{-1}\boldsymbol{Q}\right),$$
(30)

where matrix $(\sigma^2 \rho I + Q Q^T)$ is diagonal, thus its inversion is trivial. To prove that QQ^T is diagonal, recall that $Q = [Q_1 Q_2 \cdots Q_M]$, thus

$$\boldsymbol{Q}\boldsymbol{Q}^{T} = \sum_{m=1}^{M} \boldsymbol{Q}_{m} \boldsymbol{Q}_{m}^{T}.$$
(31)

The element (i, j) of matrix $Q_m Q_m^T \in \mathbb{R}^{N \times N}$ is the inner product between the *i*-th and the *j*-th rows of Q_m . Since each pixel in the *m*-th patch contributes to one and only one pixel in the synthesized image, the rows of Q_m have disjoint support, thus $(i \neq j) \Rightarrow$ $(Q_m Q_m^T)_{i,j} = 0$, that is, $Q_m Q_m^T$ has no non-zero elements outside of its main diagonal, thus is a diagonal matrix. Finally, since QQ^T is a sum of diagonal matrices, it is a diagonal matrix. It is also easy to obtain explicitly the elements of the diagonal of QQ^{T} . The diagonal elements $(Q_{m}Q_{m}^{T})_{i,i}$ are given by

$$(\boldsymbol{Q}_{m}\boldsymbol{Q}_{m}^{T})_{i,i} = \sum_{j=1}^{n} ((\boldsymbol{Q}_{m})_{i,j})^{2}.$$
 (32)

If the *m*-th patch contributes to pixel *i*, the sum in (32) contains exactly one non-zero term, thus it is equal to the square of the weight with which the *m*-th patch contributes to the synthesis of pixel *i*. If the *m*-th patch does not contributes to pixel *i*, then $(\boldsymbol{Q}_m \boldsymbol{Q}_m^T)_{i,i} = 0$. Finally, since the weight with which each patch element contributes to each pixel equals the inverse of the number of patches that contributes to that pixel, $(\boldsymbol{Q}\boldsymbol{Q}^T)_{i,i}$ is equal to the inverse of the number of patches that contribute to pixel *i*. Referring to the example in (16)–(17), we have simply $\boldsymbol{Q}\boldsymbol{Q}^T = (1/2)\boldsymbol{I}$, because each pixel is synthesized from two patches.

Finally, letting $\boldsymbol{q} = \text{diag}((\sigma^2 \rho \boldsymbol{I} + \boldsymbol{Q} \boldsymbol{Q}^T)^{-1})$, and denoting $\boldsymbol{s}^{(t)} = \boldsymbol{Q}^T \boldsymbol{y} + \sigma^2 \rho(\boldsymbol{u}^{(t)} + \boldsymbol{d}^{(t)})$, we can write the update equation (29) as

$$\boldsymbol{z}^{(t+1)} = \frac{1}{\sigma^2 \rho} \Big(\boldsymbol{s}^{(t)} - \boldsymbol{Q}^T \big(\boldsymbol{q} \odot (\boldsymbol{Q} \boldsymbol{s}^{(t)}) \big) \Big), \quad (33)$$

where \odot denotes element-wise product between two vectors. The leading cost of this update is that of the matrix-vector products involving Q and Q^T , which is O(MNn); all the other operations in (29) have lower computational cost.

7. CONCLUSIONS AND FUTURE WORK

In this paper, we have revisited patch-based image priors under the light of the synthesis vs analysis dichotomy. After showing that the classical patch-based image models (namely the EPLL) corresponds to an analysis formulation, we have proposed a patch-synthesis formulation, and analyzed its relationship with the analysis formulation, showing that they are, in general, not equivalent. Finally, we have shown how to address image denoising under the proposed formulation, via an ADMM algorithm.

We stress again that the purpose of this paper is not to introduce a new particular image prior, but a general patch-based synthesis formulation/framework, which (to the best of our knowledge) was missing from the literature on patch-based image processing, and which can be instantiated with many different patch models/priors. For this reason, we have abstained from presenting experimental results; these would critically depend on the choice and estimation of a particular patch model, which is not the focus of this paper.

Ongoing work includes the development of efficient algorithms for learning patch models under the assumption that they will be used in a synthesis formulation.

8. REFERENCES

- [1] H. Bauschke, P. Combettes, *Convex analysis and monotone* operator theory in Hilbert spaces, Springer, 2011.
- [2] S. Boyd1, N. Parikh, E. Chu, B. Peleato, J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers", *Foundations and Trends in Machine Learning*, vol. 3, pp. 1—122, 2011.
- [3] A. Buades, B. Coll, J. M. Morel, "A non-local algorithm for image denoising", Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), pp. 60-65, 2005.

- [4] Y. Chen, R. Ranftl, T. Pock, "Insights into analysis operator learning: from patch-based sparse models to higher order MRFs", *IEEE Trans. Image Proc.*, vol. 23, pp. 1060–1072, 2014.
- [5] K. Dabov, A. Foi, V. Katkovnik, K. Egiazarian, "Image denoising with blockmatching and 3D filtering", *Proc. SPIE*, vol. 6064, 2006.
- [6] M. Elad and M. Aharon, "Image denoising via sparse and redundant representations over learned dictionaries", *IEEE Trans. Image Proc.*, vol. 15, pp. 3736–3745, 2006.
- [7] M. Elad, P. Milanfar, and R. Rubinstein, "Analysis versus synthesis in signal priors", *Inv. Prob.*, vol. 23, pp. 947–968, 2007.
- [8] G. Hinton, "Training products of experts by minimizing contrastive divergence", *Neural Comp.*, vol. 14, pp. 1771–1800, 2002.
- [9] M. Lebrun, A. Buades, and J. M. Morel, "A nonlocal Bayesian image denoising algorithm", *SIAM Jour. Imaging Sciences*, vol. 6, pp. 1665–1688, 2013.
- [10] V. Papyan, M. Elad, "Multi-scale patch-based image restoration", *IEEE Trans. Image Proc.*, vol. 25, pp. 249–261, 2015.
- [11] I. Selesnick and M. Figueiredo, "Signal restoration with overcomplete wavelet transforms: comparison of analysis and synthesis priors", *Proc. SPIE*, vol. 7446 (Wavelets XIII), 2009.
- [12] J. Sulam and M. Elad, "Expected patch log likelihood with a sparse prior", in *Energy Minimization Methods in Computer Vision and Pattern Recognition*, pp. 99–111, Springer, 2015.
- [13] A. Teodoro, M. Almeida, and M. Figueiredo, "Single-frame image denoising and inpainting using Gaussian mixtures", *Proc. Inter. Conf. Pattern Recognition Applications and Methods*, pp. 283–288, 2015.
- [14] M.Wainwright and M. Jordan, "Graphical Models, Exponential Families, and Variational Inference", *Foundations and Trends* in Machine Learning, vol. 1, 2008.
- [15] D. Zoran and Y. Weiss, "From learning models of natural image patches to whole image restoration", *Proc. International Conference on Computer Vision (ICCV)*, 2011.