Separating Signal from Noise using Patch Recurrence Across Scales

Maria Zontak   Inbar Mosseri   Michal Irani
Dept. of Computer Science and Applied Mathematics
The Weizmann Institute of Science, ISRAEL

Abstract

Recurrence of small clean image patches across different scales of a natural image has been successfully used for solving ill-posed problems in clean *images (e.g., super-resolution from a single image). In this paper we show how this multi-scale property can be extended to solve ill-posed problems under noisy conditions, such as image denoising. While clean patches are obscured by severe noise in the original scale of a noisy image, noise levels drop dramatically at coarser image scales. This allows for the* unknown hidden clean patches *to "naturally emerge" in some coarser scale of the noisy image. We further show that patch recurrence across scales is strengthened when using* directional pyramids (that blur and subsample only in one direction). *Our statistical experiments show that for almost any noisy image patch (more than* 99%*), there exists a "good" clean version of itself* at the same relative image coordinates in *some coarser scale of the image. This is* a strong phenomenon *of noise-contaminated natural images, which can serve as a strong prior for separating the signal from the noise. Finally, incorporating this multi-scale prior into a simple denoising algorithm yields state-of-the-art denoising results.*

1. Introduction

The goal of natural image denoising is to recover a clean image I from its noise contaminated version I^N. This problem has been studied intensively, with considerable progress in recent years. A variety of methods have been proposed to solve this ill-posed problem, including anisotropic denoising [1], various wavelet based techniques (e.g., [13]), sparse dictionaries [5], and more. A leap improvement in denoising has been obtained by [3], who introduced a strong natural-image prior, which exploits the recurrence of small image patches internally, within a natural image. This has led to several advanced denoising algorithms, such as BM3D [4] and LSSC [9]. External statistics of natural image patches learned from large collections of clean images, has also been used as priors for denoising. These include methods like 'Fields of Experts' [14] and EPLL-GMM [17].

While the external methods benefit from a large space of clean patches, the internal methods benefit from an image-specific space of patches, which was shown by [16] to be very powerful. As shown by [16] and further analyzed by [12], external methods may suffer from noise-fitting problems. In contrast a big challenge for internal methods is to aggregate enough noisy patches which share the same underlying signal. This is particularly problematic when large patches are used (often necessary in the presence of severe noise [4, 9, 17]). In this paper we propose a new internal search space, which on one hand is *image-specific*, yet, contains *clean* patches – the space of coarse patches of the noisy image.

We present a strong multi-scale prior for solving ill-posed problems under severe noise, which is based on the recurrence of small patches *across different scales* of a natural image (where a coarser image scale is generated by *blurring* & *subsampling* the image). Most patch-based denoising methods perform denoising by exploiting patch repetitions *within* the same scale, whereas our prior suggests to use patch repetitions *across* different scales. While recurrence of patches within the original scale of the noisy images preserves their noise levels, recurrence of patches across different image scales does not: Noise levels reduce dramatically in coarser image scales, thus allowing for the 'clean' signal patches to 'naturally emerge' at those scales.

It was previously shown [7, 16] that small patches in *clean* natural images tend to recur across image scales. By "patch recurrence" we mean that a small patch (e.g., 5×5) re-appears "as is" (without any spatial transformation) in coarser versions of the image (see Fig. 1.a). This observation was successfully used for single-image Super-Resolution [7, 6]. It was further shown by [6] that due to local scale invariance in natural images, patches tend to recur at coarser image scales *at similar relative image coordinates* as in the fine scale.

We show that these natural-image properties can be exploited for "pulling out" cleaner versions of the signal from coarser scales of the noisy image (see Fig. 1.b), each patch from a different appropriate scale. It is further noted that for some patches (especially edge patches), the cross-scale recurrence property is strengthened when using *directional* pyramids (that blur and subsample only in one direction).

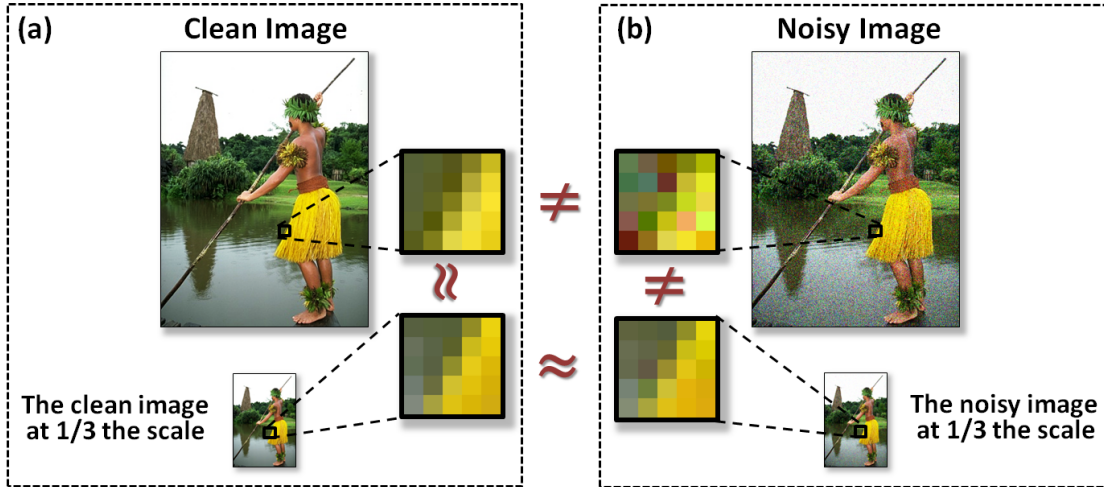Surprisingly, our experiments indicate that with very

Figure 1. *(a) A clean image and its coarser version at $1/3$ the scale. (b) A noisy image ($\sigma = 25$) and its coarser version at $1/3$ the scale. We display the 5×5 patches at the <u>same relative coordinates</u> in all 4 images (i.e., the 5×5 patch centered at (x_0, y_0) in the fine scales, and the 5×5 patch centered at $(x_0/3, y_0/3)$ in the coarse scales). There is strong similarity between corresponding 5×5 patches across-scale in the clean image. Note that the patch at the coarse scale of the noisy image is also very similar to the clean patch.*

high probability, *for every noisy patch, a very good representative of its unknown clean patch resides within a small set of 60 patches* (all its 5×5 descendant patches at the same relative image coordinates). This gives rise to a very strong prior for solving ill-posed problems under severe noise.

Finally, incorporating this multi-scale prior into a simple denoising algorithm, yields results better than state-of-the-art methods, especially for high noise levels. This is despite the fact that we use only small 5×5 patches.

Note that the term "multi-scale denoising" of [10, 2] refers to a different notion – using patches of *different sizes* from *various locations* within the image. The multi-scale property we exploit here is the "fractal-like" scale invariance property of natural images, namely, recurrence of small patches of a *fixed size* (e.g., 5×5) at the *same image location* across pyramid scales.

We believe that this prior has applications beyond denoising. It can potentially serve as a strong prior for solving other ill-posed problems in the presence of severe noise (e.g., edge-detection in noisy images, super-resolution under noisy conditions, etc.)

The rest of the paper is organized as follows. In sec. 2 we analyze and quantify the local cross-scale emergence of clean patches in noisy images, and show that this is *a very strong phenomenon*. We further show that patch recurrence across scales is strengthened when using *directional* pyramids. Sec. 3 proposes an approach for finding the "best" coarser patch for any noisy patch in a noisy image, and discusses its limitations. Sec. 4 proposes a very simple denoising algorithm which exploits this multi-scale prior, providing state-of-the-art denoising results (especially for high noise levels). Examples and results are provided in Sec. 5.

2. Statistics of Clean Patches in Noisy Images

We demonstrate how the local scale-invariance property of small patches in natural images (Fig. 1.a) can be used for "pulling out" cleaner versions of the signal from coarser scales of the noisy image (Fig. 1.b). We first extend the notion of 'pyramids' from the standard isotropic pyramid to *directional pyramids*. We then show that patch recurrence across scale is further strengthened when using directional pyramids. These ideas are illustrated in Fig. 2.

We generate a *directional pyramid* by blurring and subsampling the image only in one direction. This is different from the commonly used isotropic image pyramid, which preserves the aspect ratio, as well as from the Steerable Pyramid [15], which applies 1D directional filtering, but subsamples the image in both directions. For some patches (especially edge patches), the cross-scale recurrence property is strengthened when using directional pyramids. In particular, very thin edges which tend to disappear at low scales of the isotropic (or steerable) pyramid, are preserved well by our directional pyramids. Fig. 2 graphically illustrates the idea. The 'hidden' clean patches are obscured by noise at the original input image scale. However, these (unknown) clean patches recur at coarser image scales, at the same relative image coordinates. Since the noise level drops dramatically at coarser image scales (see Fig. 2.b), the unknown clean patches *naturally emerge* at some coarser pyramid levels. Note that although the directional pyramids change the aspect ratio of the image, the patches compared across scales remain of the same size and shape (square 5×5 patches) in all scales.

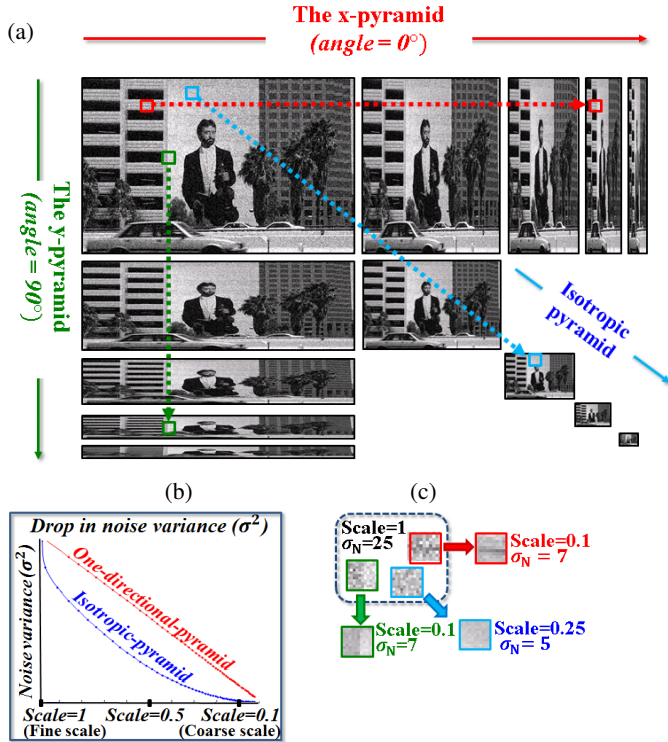We next explore how strong this scale-invariance phe-

Figure 2. **Scale invariance in "Directional Pyramids".** *(a) The 3 pyramids. (b) The drop in noise variance σ^2 as a function of the pyramid scale (empirically evaluated from many random-Gaussian-noise images). (c) Examples of noisy patches and their corresponding clean patches 'emerging' at coarse pyramid scales. Patches are marked in (a) and magnified in (c): The noisy horizontal (red) and vertical (green) edge patches have corresponding clean patches at coarse scales of the respective directional-pyramid (at the same relative coordinates). The noisy uniform patch (cyan) has a clean patch at coarse scales of the Iso-pyramid.*

nomenon is in noisy images. Namely: (i) How many 5×5 patches in a noisy image have a good clean representative patch of the clean patch in a coarser scale of the noisy image? and, (ii) How 'good' is this representative patch? (i.e., how similar to the ground-truth clean patch?)

We show that statistically this is indeed *a very strong phenomenon of noisy images*. In fact, as will be shown, the vast majority of image patches (more than 99%) will benefit significantly from going down in scale, and will have a *very good* representative patch directly below them, *at the <u>same</u> relative image coordinates* (i.e., on the 'needle' of their descendant patches – see Fig. 3).

The experiments described below were performed on hundreds of noise-contaminated natural images. Clean images (size 320x480) were taken from the *Berkeley Segmentation Database*, which contains a variety of scenes: natural, man-made, indoor, outdoor. Each clean image I (converted to grayscale) was first contaminated by Gaussian noise with zero mean and variance σ^2, resulting in a noisy image $I^N = I + N$. We then generated from I^N three

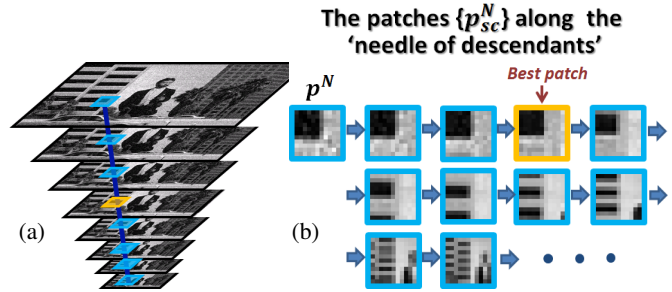The patches $\{p_{sc}^N\}$ along the 'needle of descendants'

Figure 3. **The multi-scale 'needle' of patches.** *Each noisy patch p^N has a needle in all 3 pyramids (illustrated here only for the Isotropic pyramid). (a) All the patches along the needle of the noisy patch are at the same relative image coordinates. Initially the patches get better (cleaner), but eventually new structures enter the patch. The "best" representative patch on the needle is marked in orange. (b) Zooming in on the descendant patches $\{p_{sc}^N\}$ along the needle ($sc = 1, .., 0.1$).*

types of pyramids: (i) Isotropic pyramid (blur and subsample[1] both in x and in y), (ii) X-pyramid (blur and subsample only in the x direction), and (iii) Y-pyramid (blur and subsample only in the y direction) – see Fig. 2. Each pyramid is a cascade of images $\{I_{sc}^N\}$ of gradually decreasing scales, generated by scaling down the noisy image I^N using scale factors of $sc = 0.9^s$ ($s = 0, 1, .., 20$). The smallest scale was $sc = 0.9^{20} \approx 0.1$ of the original image I^N.

Each *clean* 5×5 patch p from the clean image I was compared against the *noisy* 5×5 patches p_{sc}^N in each of the three *noisy* pyramids, but only along its 'needle of descendants' – at the <u>same</u> relative image coordinates in the coarser scales (see Fig 3). Namely, if (x, y) are the coordinate of the clean 5×5 patch p (and the noisy patch p^N), then for each scale $sc = 0.9^s$ we compare p only against the 5×5 patch p_{sc}^N whose coordinates are: (i) $(0.9^s x, 0.9^s y)$ in the Isotropic-pyramid, (ii) $(0.9^s x, y)$ in the X-pyramid, (iii) $(x, 0.9^s y)$ in the Y-pyramid. Among these descendant patches $\{p_{sc}^N\}$ ($sc = 1, .., 0.1$), we chose the one which minimizes $err_{sc} = ||p - p_{sc}^N||$ as the "best" representative of the clean patch p.

Fig. 4 visually displays the resulting image when *each patch* in the noisy image is *replaced* by its <u>single</u> "best" descendant 5×5 patch p_{sc}^N along its needle. (Overlaps between neighboring patches were simply averaged.) It is evident that the resulting image is significantly cleaner than the noisy input. In fact, it is significantly cleaner (has significantly larger PSNR[2]) than what can be achieved by today's state-of-the art denoising algorithms (see PSNR comparisons in Fig. 8). Note that this is *not a denoising algorithm*, since we used the original clean image to guide the selection of patches. However, those patches were selected from the pyramid of the *noisy* image; no additional processing was

---

[1]using Matlab "imresize" with a bicubic kernel
[2]PSNR$= 20 \log_{10}(255/\sigma)$, stands for "Peak Signal to Noise Ratio".

(a) Noisy image I^N $\Rightarrow$ "Oracle Image"
composed of *coarse patches*
of the *noisy image I^N*

(b) *Top:*   Noisy image I^N
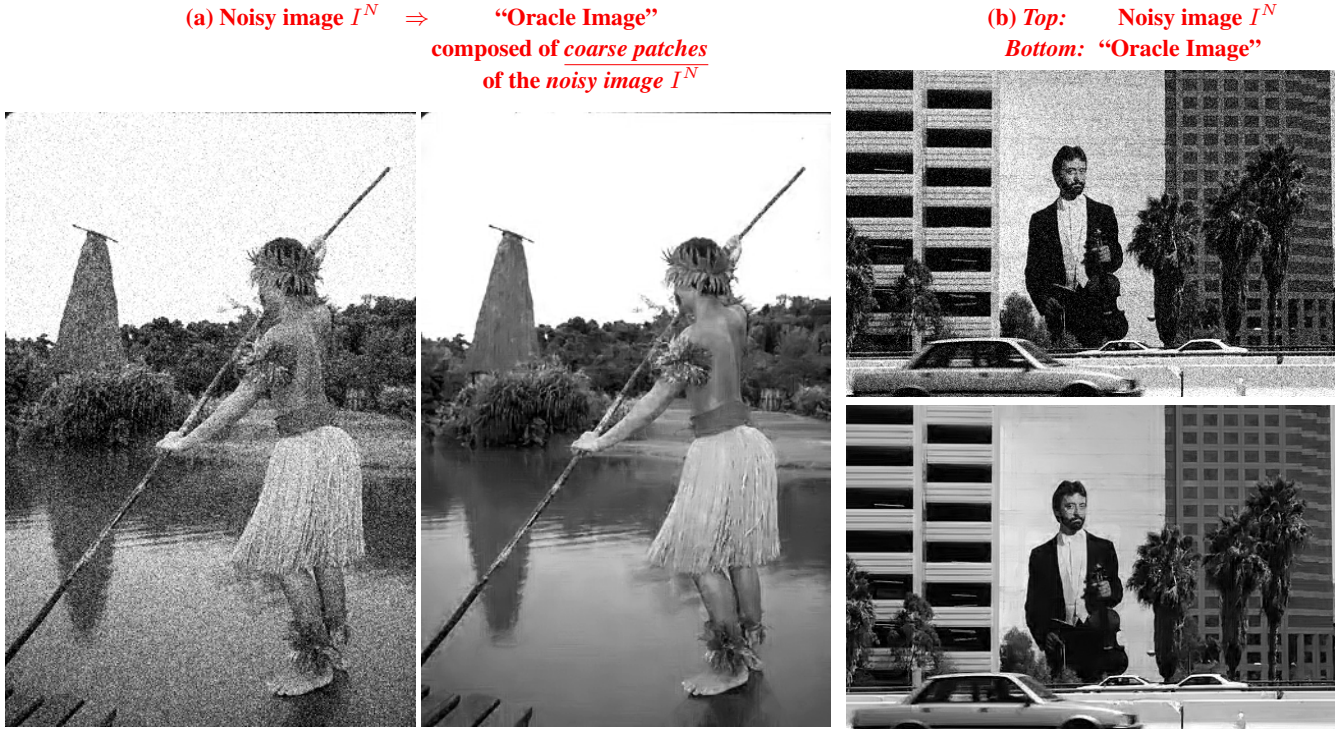*Bottom:* "Oracle Image"

Figure 4. **The "Oracle image".** *Noisy images (man-made & natural scenes), and their corresponding 'Oracle' images – generated by replacing each noisy patch with its 'best' descendant patch along its limited 'multi-scale needle' in the noisy image pyramid (see text).*

done to improve the resulting image; only a *single patch* was selected, and in a very limited search space – only along the multi-scale 'needle' descending from the noisy patch. We refer to this as the **"Oracle image"**. This clearly indicates that ***the information is there.*** *A cleaner patch exists at a coarser pyramid level of the noisy image, somewhere directly 'underneath' the noisy patch, at the same relative image coordinates.*

This is quite surprising. Note that the Oracle selection is restricted to an extremely limited search range, namely: For each noisy patch, select <u>one</u> of 60 'candidate' patches (20 patches on its needle × 3 pyramids). This is a tiny search-space relative to the *huge* theoretical space of all possible 5×5 patches (256^{25}, assuming 256 graylevels), or relative to the huge space of all natural clean image patches, or even relative to the space of all patches within the noisy image (hundreds of thousand of patches, all of which are noisy).

We ran the above Oracle experiment on a hundred natural images, repeatedly for different noise levels ($\sigma = 15, 25, 35, 45, 55, 75$). The quantitative observations reported below were averaged over all the images. Average PSNR values of the Oracle images are shown in Fig. 8, and are *well beyond the bounds of today's denoising algorithms*.

Fig. 5.a shows the mean and variance of the "best scales" (those corresponding to the "best patches") as a function of the noise level σ in I^N. As can be seen, for most noisy patches, the good representative patches tend to be in rela-

tively low scales. Moreover, these 'optimal scales' tend to further decrease as the noise level goes up. This behavior is quite intuitive, as there is a tradeoff between two factors: On one hand, we introduce more inaccuracies in the clean signal as we go down the scale (increasing 'bias'). On the other hand, the noise levels drop dramatically as we go down the scales (decreasing variance). The total patch error is the sum of these two factors: the signal error plus the noise (bias/variance tradeoff). This is illustrated in Fig. 6.b. Therefore, when there is no noise ($\sigma = 0$, $I^N = I$), it is best to stay at finest scale. However, as the noise gets larger, there is growing benefit in going down the scales.

Fig. 5.b shows the distribution of "best patches" across pyramids. Extremely few patches prefer to stay at the top scale (e.g., 0.7% for $\sigma = 25$). The majority of patches prefer the Isotropic pyramid (65% for $\sigma = 25$). These are mostly smooth patches (which are a large majority of the image). The remaining patches (mostly edges and smooth patches near the edges) prefer the directional pyramids (19% prefer the x-pyramid and 15% prefer the y-pyramid for $\sigma = 25$).

Note that diagonal edges and textured areas tend to have larger errors in the Oracle image than other parts (e.g., see the pole in the water in the lower left part of Fig. 4). This is due to using directional pyramids only in the X and Y directions. Adding directional pyramids at various other angles is bound to improve the resulting image, both visually

(a) 'Best scale' vs. σ
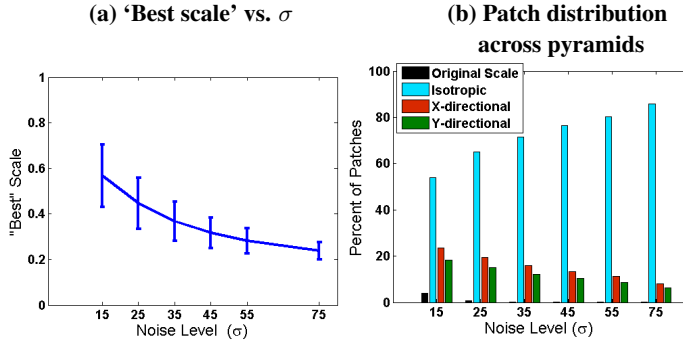

(b) Patch distribution across pyramids

Figure 5. **Quantitative observations.** *Statistics accumulated on 100 natural images, repeated with different added noise levels ($\sigma = 15, 25, 35, 45, 55, 75$).*

and numerically. In Sec. 4 we explain how the directional pyramids can be extended to *arbitrary angles*, to provide better representatives for arbitrary patch orientations. Nevertheless, most oriented patches still manage to find reasonable representatives in coarse levels of the isotropic pyramid (e.g., see the diagonal patch in Fig. 1). Hence, the above 3 pyramids alone (with their very limited search space) already provide a good representation of the clean image, with very low errors.

It is important to note that the "best" Oracle patch (the "best" patch along the the multi-scale 'needle') is *not necessarily* the best 'Nearest Neighbor' (NN) of the clean patch. In fact, it is most likely that the best NN of the clean patch resides elsewhere in the noisy pyramid. Moreover, a better representative can surely be found in a *huge* collection of clean natural patches. Nevertheless, what our experiments indicate is that there exists a *very good representative* of the clean patch in a *tiny* well-defined search space: one of 60 patches. This can give rise to a *very strong prior* for solving ill-posed problems under severe noise. In Sec. 4 we propose a somewhat naive algorithm which exploits such a prior in the context of image denoising. Although simple, this algorithm obtains state-of-the-art denoising results. Yet, it is still very far from the Oracle performance. Incorporating such a prior into more sophisticated denoising algorithms (e.g., [4, 9, 17]) is likely to lead to further improvements.

3. Estimating the Multi-Scale Patch Errors

To mimic the oracle, we wish to estimate the mean squared error (MSE) between the clean patch p and each of the noisy patches along its corresponding 'multi-scale needle' in the noisy pyramid. The problem is that we do not have the original clean patch p, only its noisy version p^N. We next show how the MSE errors with respect to the unknown clean patch p can be estimated using p^N.

Let $p^N = p + n$ be the noisy patch, where $n \sim N(0, \sigma^2)$ is assumed to be Gaussian noise (i.i.d. with zero mean and known variance σ^2). Let $p^N_{sc} = p_{sc} + n_{sc}$ denote a coarse-scale patch along the 'multi-scale needle' of p^N, where p_{sc}

denotes the coarser version of the *clean* patch p, and n_{sc} is a coarser version of the original noise n (at scale "sc"). We next show how to estimate $||p - p^N_{sc}||$ using p^N:

$$|| \underbrace{p^N}_{p+n} - p^N_{sc}||^2 = ||p - p^N_{sc}||^2 + ||n||^2 + 2 < p - p^N_{sc}, n >$$

$$\Rightarrow ||p - p^N_{sc}||^2 = ||p^N - p^N_{sc}||^2 - ||n||^2 - 2 < p - p^N_{sc}, n > \quad (1)$$

The first term on the right-hand side, $||p^N - p^N_{sc}||^2$, can be calculated from the noisy input image and its coarse-scale version. The exact noise realization in each patch is not known. Instead, we use the *expected* values of the second and third terms: $E[||n||^2] + 2E[< p - p^N_{sc}, n >]$. $n \sim N(0, \sigma^2)$, therefore $E[||n||^2] = M^2\sigma^2$, for an $M \times M$ patch. The expected value of the third term is

$$E[< p - \underbrace{p^N_{sc}}_{p_{sc} + n_{sc}}, n >] = E[< p - p_{sc}, n >] - E[< n_{sc}, n >] \quad (2)$$

The noise n is independent of the clean patch p and its coarser patches p_{sc}, therefore $E < p - p_{sc}, n >= 0$. On the other hand, $E < n_{sc}, n > \neq 0$, due to correlations across scales introduced in the downscaling process. These correlations can be either calculated empirically (averaged over hundreds of pure random noise realizations), or derived analytically (see derivation in the **www.wisdom.weizmann.ac.il/~vision/MultiScaleDenoising.html**). These noise correlations decrease as the scale gets coarser. Therefore, the error of the noisy coarse-scale patch p^N_{sc} with respect to the unknown clean patch p can be estimated by:

$$||p - p^N_{sc}||^2_{Estimate} \triangleq ||p^N - p^N_{sc}||^2 - M^2\sigma^2 + 2E[< n_{sc}, n >] \quad (3)$$

We next show how the estimate in Eq. 3 can be exploited for image denoising.

4. Using the Multi-Scale Prior for Denoising

If the error estimations in Eq. (3) were exact, then we would practically be done. All we would need to do is the following: For each noisy patch p^N, estimate the errors $||p - p^N_{sc}||^2$ for all the patches p^N_{sc} along its 'multi-scale needle' in the 3 pyramids (a total of 60 patches). Among these, the patch most similar to the (unknown) clean patch p is: $\qquad \tilde{p}^N_{sc} = arg\min_{sc} ||p - p^N_{sc}||^2_{Estimate}$
Such a procedure would yield exact 'Oracle' performance.

However, Eq. (3) is only an approximation of Eq. (1). It assumes an *average* noise behavior for all patches in the image. Thus, for example, if the empirical noise variance within a specific patch deviates from the expected variance σ^2 (as is frequently the case – see Fig. 6.a), this will result in inaccurate error estimation. E.g., if the actual (empirical) noise variance is larger than σ^2, the estimated 'optimal scale' will be higher than the true 'optimal scale' (see Fig. 6.b). To account for such deviations, we consider all patches, whose estimated error lies *within a small margin C* around Err_{min}, as good "*Best-Patch*" candidates for p^N:

5

(a) Varying patch noise levels for Gausian noise $N(0, \sigma^2)$
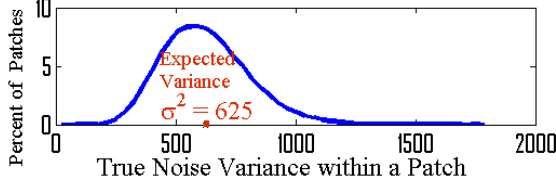
(b) Optimizing the trade-off between 2 factors:
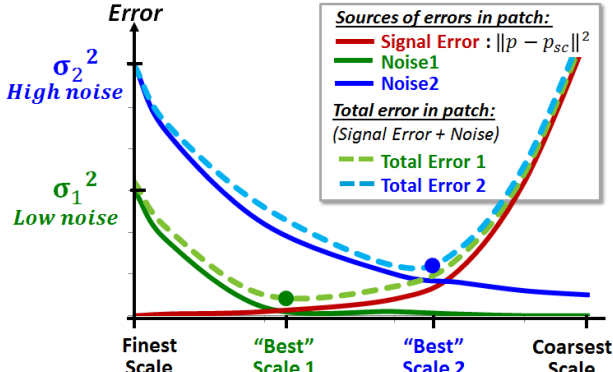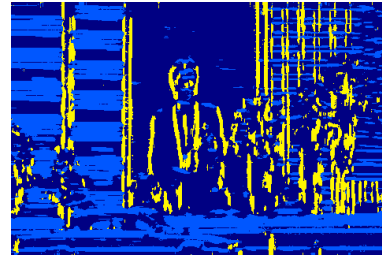**PRESERVING SIGNAL & REDUCING NOISE**

Figure 6. *Deviation in noise variance, and its effect on Optimal-Scale estimation*: (a) Noise variance in 5x5 patches often deviates from the expected σ^2. (b) Errors are introduced into the clean signal as the image scale decreases (**increasing 'bias'** - illustrated schematically; the signal degradation is patch-specific). On the other hand, noise levels drop dramatically with scale (**decreasing variance**). The total patch error is the sum of these two factors (bias/variance tradeoff). Not knowing the true noise variance within a patch may lead to inaccurate estimation of its Best-Scale.

$$G(p^N) \triangleq \{p_{sc}^N \mid ||p - p_{sc}^N||_{Estimate} < C \cdot Err_{min}\},$$

where $Err_{min} = \min_{sc} ||p - p_{sc}^N||_{Estimate}$, and $C = 1.1$ (i.e., a margin of 10% on the estimated error). ***Our algorithm boils down to simple averaging of all the 'good' representatives of p^N in the set $G(p^N)$, to obtain an estimate of its clean patch p.***

Note that the 'good' representative set $G(p^N)$ may contain patches from all pyramids. However, for oriented patches, this set will be dominated by patches from the 'correct' directional-pyramid. This is because an edge that is blurred and subsampled in the 'wrong' direction, will change its orientation with scale. This leads to very few good matches across scales in a wrongly-oriented pyramid. In contrast, the signal in an edge-patch within the correct directional pyramid is better preserved. Hence, the number of good representative patches (with small estimated error) from this pyramid will be higher, leading to automatic emphasis of this direction in the patch averaging.

Figure 7 illustrates this idea. For each 5×5 noisy patch p^N in an image with severe noise ($\sigma = 35$), we estimate its set $G(p^N)$ of 'good' representative patches. The center pixel of p^N is colored by the 'preferred' directional-

(a) Preference among 3 pyramids
(Isotropic, X-pyr and Y-pyr)   **Colorbar**
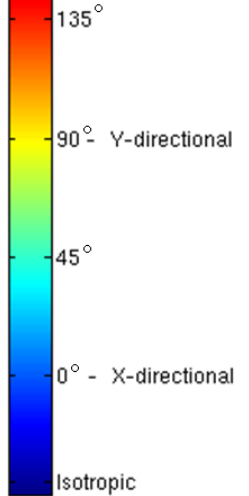
(b) Preference among 17 pyramids
(Isotropic + 16 directional pyramids)

Figure 7. **Preferred Directional-Pyramid:** *The preferred pyramid estimated by the algorithm for each 5×5 noisy patch (shown for a very noisy version of the image of Fig. 2 , with $\sigma_{noise} = 35$).*

pyramid for that patch (i.e., the pyramid that has the largest number of patch representatives in $G(p^N)$). Figure 7.a displays the result for 3 pyramids (Isotropic, X-pyr, and Y-pyr). Note that the horizontal edges (e.g., the elongated edges on the windows) prefer the X-directional pyramid, while vertical edges prefer Y-directional pyramid. Large uniform regions prefer the Isotropic pyramid. Uniform regions nearby edges benefit more from the directional pyramids.

However, oriented edges in other directions are not represented well enough by the 3 pyramids. For example, the shoulders of the man or the diagonal sides of the car, do not benefit from either the X or Y directional pyramid. These prefer the Isotropic pyramid (although not optimal, at least it does not change their orientation). We next show how directional pyramids can be extended to more directions, allowing better representation of edges in various angles.

Extension to additional directional pyramids: An edge of angle θ will obviously benefit most from blurring and subsampling in the direction θ. For any arbitrary angle θ, we can generate a θ-directional pyramid as follows: (i) Rotate the image by $-\theta$. (ii) Generate an X-pyramid from the rotated image (by blurring and subsampling in the x-direction). (iii) Rotate each pyramid level back by θ.

We extended the denoising algorithm to employ 17 pyramids: 16 directional pyramids and an Isotropic pyramid. The 16 directions were taken at equal angular gaps of $\Delta\theta = 180°/16 = 11.25°$. Although the use of 17 pyramids increases the search space of possible solutions, this does

σ_{noise}	**Ours**	BM3D [4]	LSSC [9]	EPLL [17]
15	30.93	30.87	**31.05**	30.99
25	**28.49**	28.35	**28.46**	**28.47**
35	**27.04**	26.9	26.98	26.98
45	**26.03**	25.82	25.9	25.93
55	**25.25**	25.11	25.11	25.13
75	**24.12**	24.07	23.93	23.97

Table 1. **Average PSNR values (dB)** *on 100 images with additive Gaussian noise ($\sigma = 15, 25, 35, 45, 55, 75$).*

not weaken our algorithm. On the contrary, it improves its performance. Many more edge directions now have an optimal pyramid, which preserves the edge in the respective direction. Such a pyramid also provides the largest number of 'good' candidate patches in $G(p^N)$ for denoising. Figure 7.b displays for each noisy patch its pyramid preference among 17 pyramids. The diagonal edges that earlier preferred the Isotropic pyramid, now prefer new directional pyramids in their respective directions. On average, denoising with 17 pyramids achieves an increase in PSNR of $0.1 - 0.2$ dB relative to 3 pyramids.

Removing small residual noise: The main purpose of averaging representative patches in $G(p^N)$, is to form a majority vote to the correct directional pyramid. These representative patches are already quite clean (as they come from coarse scales), but may still have *small* residual noise. The above averaging will not remove this noise, since patches from nearby scales of the same pyramid tend to have correlated noise.

However, patch similarity exists not only *across* scales, but also *within* scales. Patches in the same scale have very little noise correlation (observed and used by many denoising methods). Thus, to further improve the quality of the representative patches, we perform simple localized denoising of the representative patches of $G(p^N)$ within-their own coarse scale. While the original image scale contains severe noise, the noise level is dramatically reduced in the coarse scale. Therefore, even a simple denoising method like Non-Local Means (NLM) [3] is capable of removing *small* residual noise at such scales. Each 5×5 representative patch in $G(p^N)$ is thus cleaned using a very localized NLM (restricted to a small neighborhood of radius 3 around the patch) within its own coarse scale[3]

5. Experiments and Results

We evaluated our algorithm on 100 natural images from the Berkeley Segmentation Test Database (*BSD*), contaminated with additive Gaussian noise (with $\sigma = 15, 25, 35, 45, 55, 75$). Figure 8 and Table 1 show PSNR comparisons of several algorithms: (a) Our algorithm (17 pyramids); (b) BM3D [4]; (c) EPLL-GMM [17];

---

[3]The same neighborhood size is used for all scales and all noise levels. The σ (assumed noise level) drops as a function of the scale (Fig 2 b.).
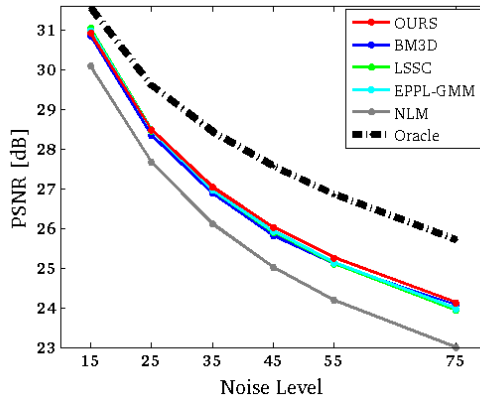

Figure 8. **Denoising results:** *Comparison of denoising algorithms on 100 noisy images (for $\sigma = 15, 25, 35, 45, 55, 75$). The black dashed line shows the 'Oracle' upper-bound. Our method exploits the multi-scale prior, providing state-of-the-art results, especially for larger noise levels. However, it is still far from the Oracle.*

(d) LSSC [9]; (e) NLM (using the implementation of $www.ipol.im/pub/art/2011/bcm\_nlm$). (b)-(d) are the current state-of the art; we used the implementations provided by the authors. Some comparative visual results can also be viewed in Fig. 9 and in the **www.wisdom.weizmann.ac.il/~vision/MultiScaleDenoising.html**.

Although simple, our algorithm achieves comparable and even slightly better results than state-of-the-art, especially at high noise levels ($\sigma \geq 35$). E.g., it achieves 0.12 db improvement over state-of-the-art for $\sigma = 55$. This is particularly surprising given that our algorithm uses only small 5x5 patches *for all noise levels*, whereas other methods need larger patches for higher noise (e.g., [4, 9] use 12x12 patches for σ=40). Moreover, [8] shows that under generic priors, *"results of BM3D are already close to optimality, and cannot be further improved beyond 0.1 dB values."* They predict poor denoising performance when using small patches in presence of severe noise. Our algorithm improves over BM3D by 0.19 dB for σ=45 (0.14 dB, σ=55), despite using 5x5 patches. This indicates the power of our local multi-scale prior (in contrast to more global priors).

Note that the initial image size limits the number of possible pyramid scales. Despite the fact that we ran our experiments on relatively small images (320×480), we obtained state-of-the art results. Applying the algorithm to larger images may lead to further improvement. Similarly, combining our multi-scale prior with more sophisticated denoising methods (e.g., [4, 9, 17]) may lead to further improvement.

Finally, although our algorithm assumes additive Gaussian noise (as do [3, 4, 17, 9]), we checked its performance on other types of noise, with encouraging results: Poisson (multiplicative noise) and Poisson-Gaussian (a mixture of multiplicative and additive noise; considered more realistic but more difficult). We used the procedure suggested by [11], which first transforms the
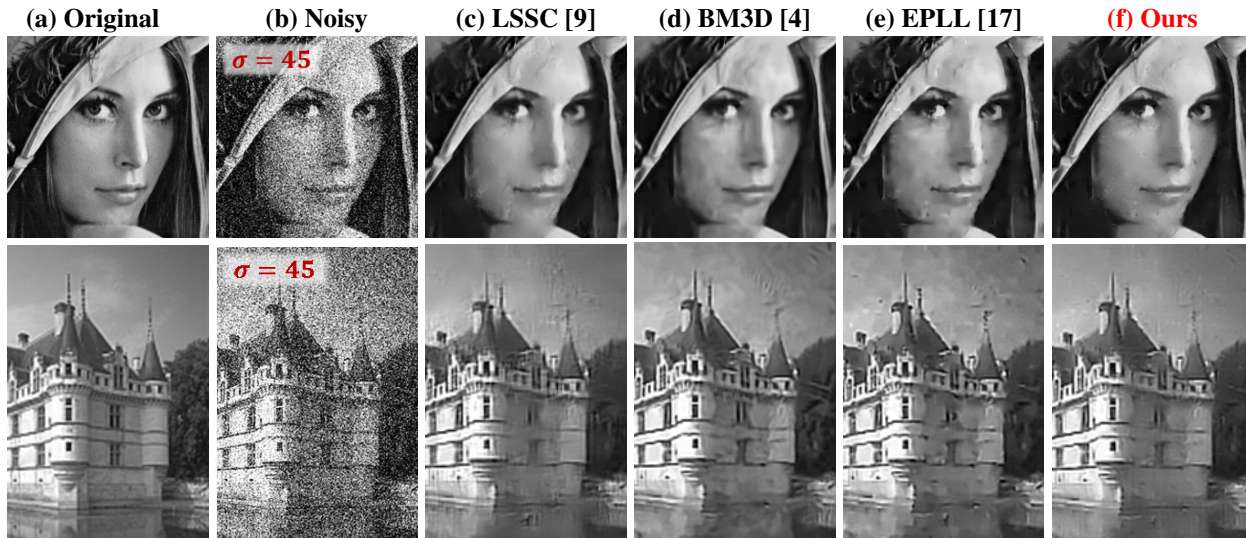
(a) Original	(b) Noisy	(c) LSSC [9]	(d) BM3D [4]	(e) EPLL [17]	(f) Ours

Figure 9. **Visual comparisons of denoising results.** *More examples -* **www.wisdom.weizmann.ac.il/~vision/MultiScaleDenoising.html**

noisy image to an additive-like domain, and applies the denoising algorithm in that domain. On the *BSD* images contaminated with Poisson-Gaussian noise, we obtained the following average PSNR results: *For peak=60, σ=6:* **27dB** (Ours), 26.84dB [4], (26.86dB) [9] and (26.92dB) [17]. *For peak=20, σ=2:* **25.76dB** (Ours), 25.63dB [4], (25.32dB) [9] and (25.54dB) [17]. The same procedure was used for all methods.

6. Summary

We present a strong multi-scale prior for solving ill-posed problems under severe noise: We observe that almost every noisy image patch has a clean version of itself hiding in some coarser scale of the image, *at the same relative image coordinates*. This restrictive search space forms a strong prior for separating the signal from the noise. Incorporating this prior into a simple denoising algorithm (simple averaging of a few coarse patches) yields state-of-the-art denoising results. This is especially true for high noise levels, despite using small 5×5 patches. This indicates the power and potential of the multi-scale prior. Combining this prior with more sophisticated denoising methods is likely to lead to further improvement. Finally, beyond denoising, we believe this can serve as a strong prior for other ill-posed problems, like edge-detection in noisy images, super-resolution in presence of noise, and more.

References

[1] L. Alvarez, P.-L. Lions, and J.-M. Morel. Image selective smoothing and edge detection by nonlinear diffusion. *SIAM J. Numer. Anal.*, 29, 1992.

[2] T. Brox, O. Kleinschmidt, and D. Cremers. Nonlocal texture filtering with efficient tree structures and invariant patch similarity measures, Aug. 2008.

[3] A. Buadess, B. Coll, and J. Morel. A non local algorithm for image denoising. In *CVPR*, 2005.

[4] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3d transform-domain collaborative filtering. *IEEE T-IP*, 16(8), 2007.

[5] M. Elad and M. Aharon. Image denoising via learned dictionaries and sparse representation. In *CVPR*, 2006.

[6] G. Freedman and R. Fattal. Image and video upscaling from local self-examples. *ACM Trans. Graph.*, 28(3), 2010.

[7] D. Glasner, S. Bagon, and M. Irani. Super-resolution from a single image. In *ICCV*, 2009.

[8] A. Levin and B. Nadler. Natural image denoising: Optimality and inherent bounds. In *CVPR*, 2011.

[9] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman. Non-local sparse models for image restoration. ICCV, 2009.

[10] J. Mairal, G. Sapiro, and M. Elad. Learning multiscale sparse representations for image and video restoration. *Multi. Model. and Simul.*, 7(1), 2008.

[11] M. Makitalo and A. Foi. Optimal inversion of the generalized anscombe transformation for poisson-gaussian noise. *IEEE T-IP*, 2012.

[12] I. Mosseri, M. Zontak, and M. Irani. Combining the power of internal and external denoising. In *ICCP*, 2013.

[13] J. Portilla, V. Strela, M. Wainwright, and E. Simoncelli. Image denoising using scale mixtures of gaussians in the wavelet domain. *IEEE T-IP*, 12(11), 2003.

[14] S. Roth and M. J. Black. Fields of experts: A framework for learning image priors. In *CVPR*, 2005.

[15] E. P. Simoncelli and W. T. Freeman. The steerable pyramid: A flexible architecture for multi-scale derivative computation. In *ICIP*, 1995.

[16] M. Zontak and M. Irani. The internal statistics of a single natural image. In *CVPR*, 2011.

[17] D. Zoran and Y. Weiss. From learning models of natural image patches to whole image restoration. In *ICCV*, 2011.