A Plurality of Sparse Representations Is Better Than the Sparsest One Alone

Michael Elad, *Senior Member, IEEE*, and Irad Yavneh

***Abstract*—Cleaning of noise from signals is a classical and long-studied problem in signal processing. Algorithms for this task necessarily rely on an *a priori* knowledge about the signal characteristics, along with information about the noise properties. For signals that admit sparse representations over a known dictionary, a commonly used denoising technique is to seek the sparsest representation that synthesizes a signal close enough to the corrupted one. As this problem is too complex in general, approximation methods, such as greedy pursuit algorithms, are often employed.**

In this line of reasoning, we are led to believe that detection of the *sparsest* representation is key in the success of the denoising goal. Does this mean that other competitive and slightly inferior sparse representations are meaningless? Suppose we are served with a group of competing sparse representations, each claiming to explain the signal differently. Can those be fused somehow to lead to a better result? Surprisingly, the answer to this question is positive; merging these representations can form a more accurate (in the mean-squared-error (MSE) sense), yet dense, estimate of the original signal even when the latter is known to be sparse.

In this paper, we demonstrate this behavior, propose a practical way to generate such a collection of representations by randomizing the Orthogonal Matching Pursuit (OMP) algorithm, and produce a clear analytical justification for the superiority of the associated Randomized OMP (RandOMP) algorithm. We show that while the maximum *a posteriori* probability (MAP) estimator aims to find and use the sparsest representation, the minimum mean-squared-error (MMSE) estimator leads to a fusion of representations to form its result. Thus, working with an appropriate mixture of candidate representations, we are surpassing the MAP and tending towards the MMSE estimate, and thereby getting a far more accurate estimation in terms of the expected ℓ_2-norm error.

***Index Terms*—Bayesian, maximum *a posteriori* probability (MAP), minimum-mean-squared error (MMSE), sparse representations.**

I. Introduction

A. Denoising in General

CLEANING of additive noise from signals is a classical and long-studied problem in signal processing. This task, known as denoising, considers a given measurement signal $\boldsymbol{y} \in \mathbb{R}^n$ obtained from the clean signal $\boldsymbol{x} \in \mathbb{R}^n$ by a contamination of the form $\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{v}$. In this paper, we shall restrict our discussion to noise vectors $\boldsymbol{v} \in \mathbb{R}^n$, assumed to be zero mean independent and identically distributed (i.i.d.) Gaussian, with entries drawn at random from the normal distribution $\mathcal{N}(0, \sigma^2)$. The denoising goal is to recover $\boldsymbol{x}$ from $\boldsymbol{y}$.

In order to design an effective denoising algorithm, we must have at our disposal two pieces of information: The first is a knowledge about the noise characteristics, as described above. Along with it, we must also introduce some knowledge about the class of signals that $\boldsymbol{x}$ belongs to. Only with these two can one design a scheme to decompose $\boldsymbol{y}$ into its original components, $\boldsymbol{x}$ and $\boldsymbol{v}$. There are numerous algorithms for denoising, as there are numerous ways to describe the *a priori* knowledge about the signal characteristics. Among these, a recently emerging model for signals that attracts much attention is one that relies on sparse and redundant representations [2], [20]. This model will be the focus of the work presented here.

B. Sparse and Redundant Representations

A signal $\boldsymbol{x}$ is said to have a sparse representation over a known dictionary $\boldsymbol{D} \in \mathbb{R}^{n \times m}$ (we typically assume that $m > n$, implying that this is a redundant representation), if there exists a sparse vector $\alpha \in \mathbb{R}^m$ such that $\boldsymbol{x} = \boldsymbol{D}\alpha$. The vector α is said to be the representation of $\boldsymbol{x}$. Referring to the columns of $\boldsymbol{D}$ as prototype signals or *atoms*, α describes how to construct $\boldsymbol{x}$ from a few such atoms by a linear combination. The representation is sparse—the number of non-zeros in it, $k = \|\alpha\|_0$, is expected to be much smaller than n. Also, this is a redundant representation—it is longer than the original signal it represents. In this paper, we consider the family of signals that admit sparse representations over a known dictionary $\boldsymbol{D}$ and discuss ways to denoise them. Note that at this stage we do not provide a full and exact definition of this signal family (e.g., we do not specify how the representations are generated)—such a definition will follow at a later stage in the paper, where a rigorous analysis is pursued.

Assuming that $\boldsymbol{x} = \boldsymbol{D}\alpha$ with a sparse representation α, how can we denoise a corrupted version of it, $\boldsymbol{y}$? A commonly used denoising technique is to seek the sparsest representation that synthesizes a signal close enough to the corrupted one [2], [10]–[14], [17], [18], [21], [28], [29]. Put formally, one way to define our task is given by

$$\hat{\alpha} = \arg\min_{\alpha} \ \|\alpha\|_0 + \lambda \|\boldsymbol{y} - \boldsymbol{D}\alpha\|_2^2. \tag{1}$$

The first penalty directs the minimization task towards the sparsest possible representation, exploiting our *a priori* knowledge about the formation of the signal. The second penalty manifests our knowledge about the noise being white and Gaussian. This overall expression is inversely proportional to

Manuscript received April 10, 2008; revised April 03, 2009. Current version published September 23, 2009. This work was supported by the Center for Security Science and Technology–Technion, and by The EU FET-OPEN project SMALL under Grant 225913.

The authors are with the Department of Computer Science, Technion–Israel Institute of Technology, Technion City, Haifa 32000, Israel (e-mail: elad@cs.technion.ac.il; irad@cs.technion.ac.il).

Communicated by L. Tong, Associate Editor for Detection and Estimation.

Color versions of Figures 3–8 in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIT.2009.2027565

the posterior probability, $p(\alpha|\boldsymbol{y})$, and as such, its minimization forms the maximum *a posteriori* probability (MAP) estimate [2]. The parameter λ should be chosen based on σ and the fine details that model how the signals' representations are generated. As remarked above, there are other ways to formulate our goal—for example, we could replace one of the penalties with a constraint, if their size is known. Once $\hat{\alpha}$ is found, the denoising result is obtained by $\hat{\boldsymbol{x}} = \boldsymbol{D}\hat{\alpha}$.

The problem posed in (1) is too complex in general, requiring a combinatorial search that explores all possible sparse supports [22]. Approximation methods are therefore often employed, with the understanding that their result may deviate from the true solution. One such approximation technique is the Orthogonal Matching Pursuit (OMP), a greedy algorithm that accumulates one atom at a time in forming $\hat{\alpha}$, aiming at each step to minimize the representation error $\|\boldsymbol{y} - \boldsymbol{D}\alpha\|_2^2$ [2], [4], [6], [7], [21], [23], [27]. When this error falls below some predetermined threshold, or when the number of atoms reaches a destination value, this process stops. While crude, this technique works very fast and can guarantee near-optimal results in some cases.

How good is the denoising obtained by the above approach? Past work provides some preliminary, both theoretical and empirical, answers to this and related questions [2], [9]–[11], [13], [14], [17], [18], [28], [29]. Most of this work concentrates on the accuracy with which one can approximate the true representation (rather than the signal itself), adopting a worst case point of view. Indeed, the only work that targets the theoretical question of denoising performance head-on is reported in [13], [14], providing asymptotic assessments of the denoising performance for very low and very high noise powers, assuming that the original combinatorial problem can be solved exactly.

C. This Paper

In the above line of reasoning, we are led to believe that detection of the sparsest representation is key in the success of the denoising goal. Does this mean that other, competitive yet slightly inferior, sparse representations are meaningless? This question is critical, especially due to the often encountered cases where OMP (and other approximation algorithms) fails to find the truly sparsest representation.

Furthermore, past analysis of approximation algorithms indicates that a measure of coherence of the dictionary can predict the tendency of these methods to fail, and this is especially true when the sparsity of the target representation is rather mild [2], [10], [11], [17], [18], [27], [28]. The coherence of a dictionary is defined through the worst pair of atoms exhibiting maximal correlation. If this value is high, it implies that these atoms tend to confuse and mislead any solver. Thus, noncoherent dictionaries necessarily lead to wrong solutions in many cases; are these to be considered as complete failures?

In fact, we should ask a more general question: Suppose we are served with a group of competing sparse representations, each claiming to explain the signal differently. Can those be fused somehow to lead to a better result? Surprisingly, the answer to this question is positive; these representations can definitely be merged to form a more accurate estimate of the original signal. This means that even when the dictionary is noncoherent, one can obtain a reasonable denoising, by exploiting this mixture of representations. Why is this true? How can we exploit this? In this paper, we aim to show that there is life beyond the sparsest representation. More specifically:

- We propose a practical way to generate a set of sparse representations for a given signal by randomizing the OMP algorithm. This technique samples from the set of sparse solutions that approximate $\boldsymbol{D}\alpha = \boldsymbol{y}$.
- We demonstrate the gain in using such a set of representations through a preliminary experiment that fuses these results by a plain averaging; and most important of all,
- we provide a clear explanation for the origin of this strange phenomenon. We develop analytical expressions for the MAP and the minimum mean-squared-error (MMSE) estimators for the model discussed, and show that while the MAP estimator aims to find and use the sparsest representation, the MMSE estimator requires a fusion of a collection of representations to form its result. Thus, working with a set of candidate representations, we are surpassing the MAP and tending towards the MMSE estimate, and thereby getting a more accurate estimation in the average ℓ_2 error sense.
- Based on the above rigorous analysis, we also provide clear expressions that predict the mean-square error (MSE) of the various estimators, and thus obtain a good prediction for the denoising performance of the OMP and its randomized version.

The idea to migrate from the MAP to the MMSE has been recently studied elsewhere [19], [25], [26]. Our work differs from these earlier contributions in its motivation, in the objective it sets, and in the eventual numerical approximation algorithm proposed. Our work's origin is different as it emerges from the observation that several different sparse representations convey an interesting story on the signal to be recovered, and this is reflected in the structure of this paper. The MMSE estimator we develop targets the signal $\boldsymbol{x}$, while the above works focus on estimation of the representation α. Finally, and most importantly, our randomized greedy algorithm is very different from the deterministic techniques explored in [19], [25], [26].

This paper is organized as follows. In Section II, we build a case for the use of several sparse representations, leaning on intuition and some preliminary experiments that suggests that this idea is worth a closer look. Section III contains the analytic part of this paper, which develops the MAP and the MMSE exact estimators and their expected errors, showing how they relate to the use of several representations. We conclude in Section IV by highlighting the main contribution of this paper, and drawing attention to important open questions to which our analysis points.

II. The Key Idea—A Mixture of Representations

In this section, we build a case for the use of several sparse representations. First, we motivate this by drawing intuition from example-based modeling, where several approximations of the corrupted data are used to denoise it. Armed with the desire to generate a set of sparse representations, we present the Randomized OMP (RandOMP) algorithm that generates a group of competitive representations for a given signal. Finally, we show that this concept works quite well in practice and

Task: Approximate the solution of $\min_\alpha \ \|\alpha\|_0 \ \text{ s.t. } \ \|\mathbf{D}\alpha - \mathbf{y}\|_2 \le T$.

Parameters: The matrix $\mathbf{D}$, the signal $\mathbf{y}$, and the error threshold T.

Initialization: Initialize $k = 0$, and set

- The initial solution $\alpha^0 = 0$.
- The initial residual $\mathbf{r}^0 = \mathbf{y} - \mathbf{D}\alpha^0 = \mathbf{y}$.
- The initial solution support $\mathcal{S}^0 = Support\{\alpha^0\} = \emptyset$.

Main Iteration: Increment k by 1 and perform the following steps:

- **Sweep:** Compute the errors $\epsilon(j) = \min_{z_j} \|\mathbf{d}_j z_j - \mathbf{r}^{k-1}\|_2^2$ for all j using the optimal choice $z_j^* = \mathbf{d}_j^T \mathbf{r}^{k-1} / \|\mathbf{d}_j\|_2^2$.
- **Update Support:** Find j_0–the minimizer of $\epsilon(j)$, and update the support, $\mathcal{S}^k = \mathcal{S}^{k-1} \cup \{j_0\}$.
- **Update Solution:** Compute α^k, the minimizer of $\|\mathbf{D}\alpha - \mathbf{y}\|_2^2$ subject to $Support\{\alpha\} = \mathcal{S}^k$.
- **Update Residual:** Compute $\mathbf{r}^k = \mathbf{y} - \mathbf{D}\alpha^k$.
- **Stopping Rule:** If $\|\mathbf{r}^k\|_2 < T$, stop. Otherwise, apply another iteration.

Output: The proposed solution is α^k obtained after k iterations.

Fig. 1. The OMP— a greedy algorithm.

provides a benefit over the use of a single representation. In the next section, we provide a more rigorous explanation of this phenomenon.

A. Why Bother? Some Hand-Waving

Why should we consider the use of more than just one representation? One possible reason is the fact that a single representation we happen to find may be the wrong one, due to the tendency of pursuit techniques to fail from time to time. This suggests that one should possibly seek several representations and choose the best one in some sense, assuming that this is possible. In doing so, we are still in the realm of the MAP estimator, aiming to improve the OMP chances to approximate better the desired single-representation solution of (1). As we shall see next, we aim to achieve far more than just that when using a set of representations.

A second and more relevant reason to consider a mixture of representations is an intriguing relation between our model and direct example-based techniques. Our model assumes that signals in the family we handle can be represented as multiplications of the dictionary $\boldsymbol{D}$ by sparse vectors. What if we allow $\boldsymbol{D}$ to be more redundant by letting the number of its columns m to grow? The general tendency we expect to see is a decrease in the number of non-zeros required in the representations, that is, they become sparser. At the extreme, when the dictionary contains $m \to \infty$ columns, reflecting many possible instances of signals, the required sparsity should tend towards $k = 1$, since almost every original signal is available as an atom (possibly up to a scale).

This extreme case is exactly the one practiced in direct example-based methods [1], [8], [15], [16], [24]. Suppose we are given many instances of noisy signals $\{\boldsymbol{y}_i\}_{i=1}^N$. We refer to those as our training data, and form a dictionary $\boldsymbol{D}$ by simply concatenating them as our atoms. When aiming to denoise a newly obtained signal $\boldsymbol{y}$, an example-based denoising algorithm suggests that we seek in $\boldsymbol{D}$ a set of nearby atoms. Each such neighbor found is an extremely sparse representation with cardinality 1, and with the coefficient being 1 as well. We may consider a slightly more general search for neighbors that allows for scaling, which enriches the span of the dictionary signal set.

Using one neighbor atom only as our suggested solution implies that we replace the noise in $\boldsymbol{y}$ by the noise in this atom, rather than cleaning it, which is, of course, useless. Suppose that in the set of neighbors chosen we have managed to find instances of the same original signal with different realizations of noise. In such a case, averaging these solutions leads to an attenuation of the additive noise. Thus, a collection of very sparse representations joins forces to produce a better estimate. Indeed, a closer look at the Non-Local-Means (NLM) algorithm [3] reveals that this is the very same method employed, where the dictionary used consists of patch examples extracted from the neighborhood of the sample to be cleaned.

If the above is true for the extreme case, why should it not be relevant for the lower redundancy case as well? The rationale is that each sparse representation found recommends its own way of denoising, and their fusion may lead to a better overall noise removal effect. Could this be true? In order to explore this idea, we must start by finding a practical way to generate a set of candidate representations, which is our next topic.

B. RandOMP

Here is a clear definition of our goal: Given a dictionary $\boldsymbol{D}$ and a signal $\boldsymbol{y}$, we aim to find a group of sparse representations α_i, such that each satisfies $\|\boldsymbol{D}\alpha_i - \boldsymbol{y}\|_2 \le T$, and all aim to be as sparse as possible yet different from each other. Alternatively, we may desire to find this set such that each has the same prespecified number of non-zeros, k, and all aim to get residuals, $\|\boldsymbol{D}\alpha_i - \boldsymbol{y}\|_2$, that are as low as possible. We shall work in this subsection with the former option, since it is more relevant to denoising in cases when the noise power is fixed and known, as in the case studied here.

Fig. 1 presents the OMP algorithm with a stopping rule that depends on the residual energy [2], [4], [6], [7], [21], [23]. At

Task: Approximate the solution of $\min_\alpha \ \|\alpha\|_0 \ \text{s.t.} \ \|\mathbf{D}\alpha - \mathbf{y}\|_2 \le T$.

Parameters: The matrix $\mathbf{D}$, the signal $\mathbf{y}$, and the error threshold T.

Initialization: Initialize $k = 0$, and set

- The initial solution $\alpha^0 = 0$.
- The initial residual $\mathbf{r}^0 = \mathbf{y} - \mathbf{D}\alpha^0 = \mathbf{y}$.
- The initial solution support $\mathcal{S}^0 = Support\{\alpha^0\} = \emptyset$.

Main Iteration: Increment k by 1 and perform the following steps:

- **Sweep:** Compute the errors $\epsilon(j) = \min_{z_j} \|\mathbf{d}_j z_j - \mathbf{r}^{k-1}\|_2^2$ for all j using the optimal choice $z_j^* = \mathbf{d}_j^T \mathbf{r}^{k-1}/\|\mathbf{d}_j\|_2^2$.
- **Update Support:** Draw j_0 at random with probability proportional to $\exp\{\frac{c^2}{2\sigma^2} \cdot |\mathbf{d}_j^T \mathbf{r}^{k-1}|^2/\|\mathbf{d}_j\|_2^2\}$, and update the support, $\mathcal{S}^k = \mathcal{S}^{k-1} \cup \{j_0\}$.
- **Update Solution:** Compute α^k, the minimizer of $\|\mathbf{D}\alpha - \mathbf{y}\|_2^2$ subject to $Support\{\alpha\} = \mathcal{S}^k$.
- **Update Residual:** Compute $\mathbf{r}^k = \mathbf{y} - \mathbf{D}\alpha^k$.
- **Stopping Rule:** If $\|\mathbf{r}^k\|_2 < T$, stop. Otherwise, apply another iteration.

Output: The proposed solution is α^k obtained after k iterations.

Fig. 2. RandOMP—generating random sparse representations.

each iteration, the set $\{\epsilon(j)\}_{j=1}^m$ is computed, whose jth term indicates the error that would remain if atom j were added to the current solution. The atom chosen is the one yielding the smallest error. Note that if there are several candidate atoms that show a relatively small residual energy, the smallest one is chosen regardless of the proximity of the others to it. This brings us naturally to the randomization approach we intend to apply.

In order to use this algorithm to generate a set of (probably) distinct sparse representations, all that we need to do is to randomize the choice of the next atom to be added. For example, rather than choose the atom that minimizes $\epsilon(j)$, we can choose it at random with a probability inversely proportional to these error values, or proportional to $|\boldsymbol{d}_j^T \boldsymbol{r}^{k-1}|^2/\|\boldsymbol{d}_j\|_2^2$ (since $\epsilon(j) = \|\boldsymbol{r}^{k-1}\|_2^2 - |\boldsymbol{d}_j^T \boldsymbol{r}^{k-1}|^2/\|\boldsymbol{d}_j\|_2^2$). For reasons to be explained in detail in the next section, the specific way we choose to draw the next atom is with probability linearly proportional to $\exp\{\frac{c^2}{2\sigma^2} \cdot |\boldsymbol{d}_j^T \boldsymbol{r}^{k-1}|^2/\|\boldsymbol{d}_j\|_2^2\}$, with

$$c^2 = \frac{\sigma_x^2}{\sigma_x^2 + \sigma^2}\,. \tag{2}$$

Here σ_x is the variance of the nonzero entries of the representation of the original signal.

By running this algorithm J_0 times, this randomization leads to J_0 solutions $\{\alpha_i\}_{i=1}^{J_0}$, as desired. Common to all these representations are the facts that (i) their representation error $\|\boldsymbol{D}\alpha_i - \boldsymbol{y}\|_2$ is below T due to the stopping rule enforced; and (ii) all of them tend to be relatively sparse due to the greedy nature of this algorithm that aims to decrease the residual energy, giving preference to those atoms that serve this goal better. Fig. 2 presents this algorithm.

We demonstrate the behavior of this algorithm by performing the following simple test. First, we build a random dictionary $\boldsymbol{D}$ of size 100×200 by drawing its entries at random from the normal distribution $\mathcal{N}(0,1)$, and then ℓ_2 normalizing its columns. We then generate a random representation α_0 with $k = 10$ non-zeros chosen independently at random and with values drawn from $\mathcal{N}(0, \sigma_x^2)$ with $\sigma_x = 1$. The clean signal is obtained by $\boldsymbol{x} = \boldsymbol{D}\alpha$, and its noisy version $\boldsymbol{y}$ is obtained by adding white Gaussian noise with entries drawn from $\mathcal{N}(0, \sigma^2)$ with $\sigma = 1$ as well.

Armed with the dictionary $\boldsymbol{D}$, the corrupted signal $\boldsymbol{y}$ and the noise threshold $T = n\sigma^2 = 100$, we first run the plain OMP, and obtain a representation α^{OMP} with cardinality $k = 2$, and with a representation error $\|\boldsymbol{D}\alpha^{\mathrm{OMP}} - \boldsymbol{y}\| = 99.82$. We can also check the denoising effect obtained by evaluating the expression $\|\boldsymbol{D}\alpha^{\mathrm{OMP}} - \boldsymbol{x}\|_2^2/\|\boldsymbol{y} - \boldsymbol{x}\|_2^2$. The value obtained is 0.168, suggesting that the noise was indeed attenuated nicely by a factor close to 6.

We proceed by running RandOMP $J_0 = 1000$ times, obtaining 1000 candidate representations $\{\alpha_j^{\mathrm{RandOMP}}\}_{j=1}^{1000}$. Among these, there are 999 distinct ones, but we allow repetitions. Fig. 3(a) presents a histogram of the cardinalities of the results. As can be seen, all the representations obtained are relatively sparse, with cardinalities in the range $[2, 21]$, indicating that the OMP representation is the sparsest. Fig. 3(b) presents a histogram of the representation errors of the results obtained. As can be seen, all the representations give an error slightly smaller than the threshold chosen, $T = 100$.

We also assess the denoising performance of *each* of these representations as done above for the OMP result. Fig. 3(c) shows a histogram of the denoising factor

$$\|\boldsymbol{D}\alpha_j^{\mathrm{RandOMP}} - \boldsymbol{x}\|_2^2/\|\boldsymbol{y} - \boldsymbol{x}\|_2^2.$$

The results are in the range $[0.128, 0.296]$, with roughly half surpassing the OMP performance and the other half performing more poorly. However, can we detect the better performing representations? Fig. 3(d) shows the relation between the representations' cardinality and their expected performance, and as can be seen, it is hard to choose the best performing one judging only by their cardinalities. This brings us to the next discussion about a way to fuse the results to get an enhanced overall denoising performance.

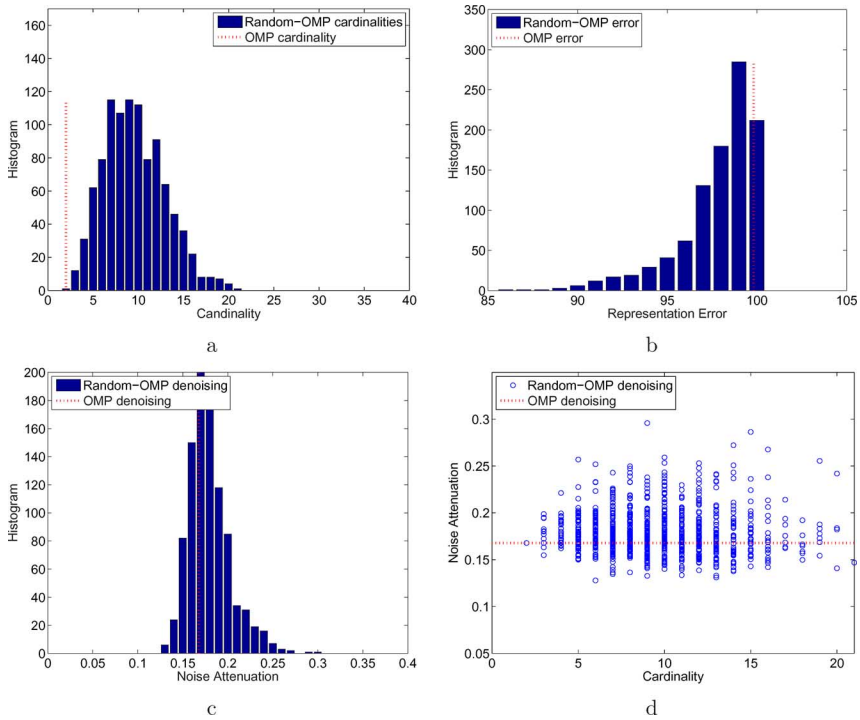

Fig. 3. Results of the RandOMP algorithm with 1000 runs. (a) A histogram of the representations' cardinalities. (b) A histogram of the representations' errors; (c) A histogram of the representations' denoising factors. (d) The denoising performance versus the cardinality.

C. Rule of Fusion

While it is hard to pinpoint the representation that performs best among those created by the RandOMP, their averaging is quite easy to obtain. The questions to be asked are as follows. (i) What weights to use when averaging the various results? (ii) Will this lead to better overall denoising? We shall answer these questions intuitively and experimentally below. In Section III, we revisit these questions and provide a justification for the choices made.

From an intuitive point of view, we might consider an averaging that gives a precedence to sparser representations. However, our experiments indicate that a plain averaging works even better. Thus, we use the formula[1]

$$\alpha^{\mathrm{AVE}} = \frac{1}{J_0}\sum_{j=1}^{J_0} \alpha_j^{\mathrm{RandOMP}}. \tag{3}$$

We return to the experiment described in the previous subsection, and use its core to explore the effect of the averaging described above. We perform 1000 different experiments that share the same dictionary but generate different signals α, $\boldsymbol{x}$, and $\boldsymbol{y}$ using the same parameters ($\sigma_x = \sigma = 1$ and $k = 10$). For each experiment, we generate $J_0 = 100$ RandOMP representations and average them using (3).

Fig. 4 presents the results—for each experiment a point is positioned at the denoising performance of the OMP and the corresponding averaged RandOMP. As can be seen, the general tendency suggests much better results with the RandOMP. The average denoising performance over all these experiments is 0.186 for the OMP and 0.105 for the averaged RandOMP method. The mean denoising factor of OMP versus that of RandOMP is denoted by a square mark. Note the concentration of some of the results on the diagonal line. This indicates that in some experiments, the OMP and the RandOMP give the same results. This takes place in cases where $\|\boldsymbol{y}\|_2 \leq T$, implying that the zero solution is found sufficient and both algorithms stop immediately.

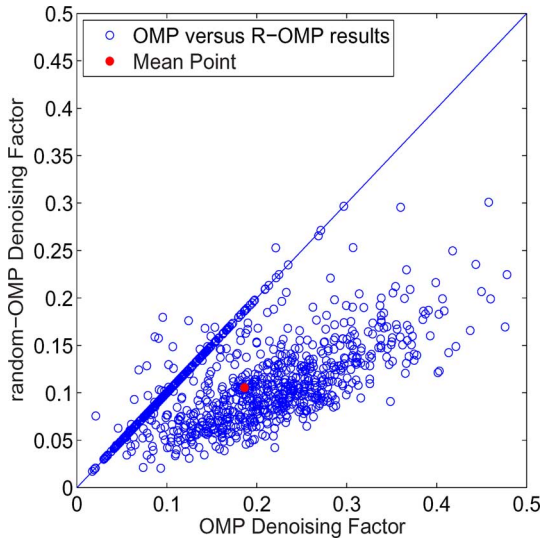

Fig. 4. Results of 1000 experiments showing the plain OMP denoising performance versus those obtained by the averaged RandOMP with $J_0 = 100$ candidate results.

The above results are encouraging and immediately lead to more questions. (i) How many different representations are

[1]In Section III, we show that both the OMP and RandOMP solutions should actually be multiplied by a shrinkage factor, c^2, defined in (2), which is omitted in this experiment.

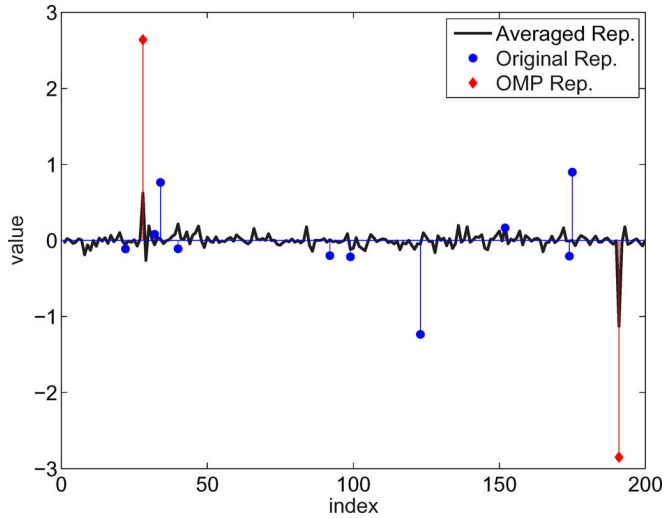

Fig. 5. The true (original) representation, the one found by the OMP, and the one obtained by averaging 1000 representations created by RandOMP.

enough in order to enjoy the gain of the RandOMP averaging? (ii) How does this gain behave as a function of the input signal-to-noise ratio (SNR)? (iii) How does this gain behave for different cardinalities of the original representation? (iv) What is the effect of the dictionary (and its redundancy) on these results? (v) Are these results related to some sort of known estimator? and most important of all, (vi) Why do we get this gain at all? We shall provide experimental answers to questions (i)-(iv) in the next subsection, and treat questions (v) and (vi) in Section III by providing a detailed analysis of the estimation problem at hand.

Just before leaving this section, we would like to draw attention to the following interesting behavior. When averaging the representations in forming the denoising estimate, we obtain a new representation α^{AVE} that is no longer sparse. Nevertheless, this representation is the one that leads to the improved results. Fig. 5 shows the true representation, the OMP one, and α^{AVE} obtained with 1000 runs of the RandOMP, in a sequel to the experiment shown in Section II-B. As can be seen, these three representations are quite distinct, and yet they lead to very similar signals (the denoising factor obtained in this case is 0.168 for the OMP, and 0.06 for the averaged representation). While the OMP uses less atoms than the original one, the averaged representation is dense, using *all* the atoms with appropriate weights.

D. Hey, It Works! Some Experiments and Results

In this subsection, we shall empirically answer some of the questions raised above, with an aim to better map the behavior of the RandOMP averaging method in different scenarios for various settings.

First we address the question of how many different representations to use in order to enjoy the gain of RandOMP averaging. As the complexity of the new estimator with J_0 different representations is about J_0 times higher than that of the plain OMP, there is a strong incentive to reduce J_0 as much as possible without sacrificing performance. Fig. 6(a) presents the averaged results over 1000 experiments, for a varying number of representations in the range $J_0 \in [5, 200]$. We see that while more representations improve the results, the lion's share of the improvement over the OMP is obtained even for small values of J_0.

All the tests done so far assumed $\sigma = \sigma_x = 1$ with $k = 10$. This case corresponds to a very low SNR of k/n and below, since the noise power is $n\sigma^2$, while the signal power is below $k\sigma_x^2$ (depending on the k atoms chosen and their relative orientations.) Thus, we must ask—how is the gain observed affected by the input SNR? In order to explore this, we fix the parameters $\sigma_x = 1$, $k = 10$, $J_0 = 40$, vary the noise power in the range[2] $\sigma \in [0.1, 2]$, and average the denoising results over 200 experiments. Fig. 6(b) presents the denoising performance of the averaging as a function of σ, and as can be seen, our method is better for all the choices of σ, but the gain it provides is higher for lower SNR—as σ grows (and SNR drops), the ratio between the two curves in Fig. 6(b) increases.

The next test we perform considers the complexity of the original signal, by varying k in the range $[1, 40]$. The sparser the representation of the original signal, the easier it is supposed to be denoised. Naturally, we desire to find out how the gain of the RandOMP average behaves for different cardinalities of the original representation. Fig. 6(c) presents the results obtained for $\sigma = \sigma_x = 1$, showing that the OMP is inferior to the averaged results for all cardinalities.

The last test we present studies the effect of the redundancy of the dictionary on the denoising performance. We fix the parameters $\sigma = \sigma_x = 1$, $k = 10$, $J_0 = 40$, the dimension of the signal is set to $n = 100$, and we vary the number of the atoms in the range $m \in [10, 400]$. Averaging the denoising results over 200 experiments, we obtain the results as shown in Fig. 6(d). These clearly show that for a wide range of redundancies, the gain obtained by the averaging of the RandOMP results remains unchanged, and the denoising factor appears to be independent of m (as opposed to the one obtained by the OMP which deteriorates). The case of underdetermined dictionaries ($m \leq n = 100$) and especially for $m \to k$ is special, since there the representations found tend to be full, leading to a convergence of the two methods (OMP and RandOMP).

We add that a similar test done on a redundant DCT dictionary[3] led to very similar results, suggesting that the behavior we observe is robust with respect to the dictionary properties.

E. Summary

When we have at our disposal several competing sparse representations of the same noisy signal, they can be averaged to provide better denoising performance. The combined representation is no longer sparse, but this does not reduce its efficiency in attenuating the noise in the signal. In this subsection, we described how to obtain such a group of representations, how to fuse them, and what to expect. Specifically, we found that the

[2]This is a wide range of SNR that provides a good coverage of the cases to be tested.

[3]This dictionary is obtained by assigning

$$d[i,j] = \cos((i-1)(j-1)\pi/m), \quad \text{for } 1 \leq i \leq n \text{ and } 1 \leq j \leq m$$

removing the mean from all the atoms apart from the first, and normalizing each atom to unit ℓ_2-norm.

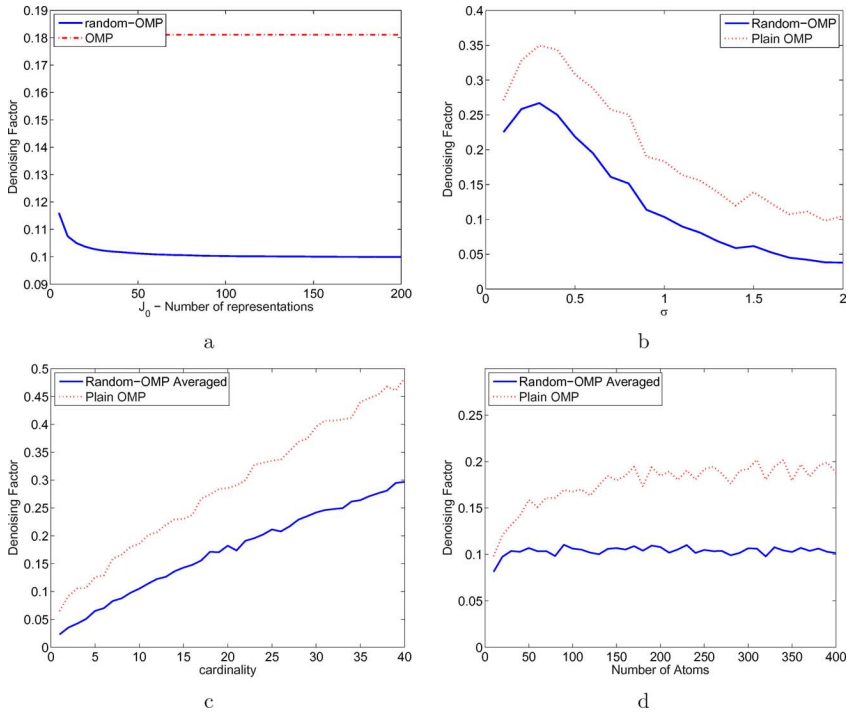

Fig. 6. Various tests on the RandOMP algorithm, checking how the denoising is affected by (a) the number of representations averaged; (b) the input noise power; (c) the original representation's cardinality; and (d) the dictionary's redundancy.

method we propose appears to be very effective, robust with respect to the signal complexity, dictionary type, and redundancy, and yields benefits even when we merge only a few representations. We now turn to provide a deeper explanation of these results by a careful modeling of the estimation problem and development of MAP and MMSE estimators.

III. Why Does it Work? A Rigorous Analysis

In this section, we start by modeling the signal source in a complete manner, define the denoising goal in terms of the MSE, and derive several estimators for it. We start with a very general setting of the problem, and then narrow it down to the case discussed above on sparse representations. Our main goal in this section is to show that the MMSE estimator can be written as a weighted averaging of various sparse representations, which explains the results of the previous section. Beyond this, the analysis derives exact expressions for the MSE for various estimators, enabling us to assess analytically their behavior and relative performance, and to explain results that were obtained empirically in Section II. Towards the end of this section we tie the empirical and the theoretical parts of this work—we again perform simulations and show how the actual denoising results obtained by OMP and RandOMP compare to the analytic expressions developed here.

A. A General Setting

1) Notation: We denote continuous (resp., discrete) vector random variables by small (resp., capital) letters. The probability density function (pdf) of a continuous random variable a over a domain Ω_a is denoted $p(a)$, and the probability of a discrete random variable A by $P(A)$. If $b_1,\ldots,b_k$ is a set of continuous (and/or discrete) random variables, then $p(a|b_1,\ldots,b_k)$ denotes the conditional pdf of a subject to $b_1 \,\&\, b_2, \ldots \,\&\, b_k$. Similarly, $P(A|B_1,\ldots,B_k)$ denotes the conditional probability for a discrete event A. With $\mathcal{E}$ denoting expectation, we denote the mean of a by

$$\mathcal{M}(a) = \mathcal{E}(a) = \int_{\Omega_a} a p(a)\, da$$

and the variance by

$$\mathcal{V}(a) = \mathcal{E}\left(\|a - \mathcal{M}(a)\|^2\right) = \int_{\Omega_a} \|a - \mathcal{M}(a)\|^2 p(a)\, da.$$

Similarly, in the discrete case

$$\mathcal{M}(A) = \mathcal{E}(A) = \sum_{\Omega_A} A P(A)$$

and

$$\mathcal{V}(A) = \mathcal{E}\left(\|A - \mathcal{M}(A)\|^2\right) = \sum_{\Omega_A} \|A - \mathcal{M}(a)\|^2 P(A).$$

Finally, we denote conditional means and variances by

$$\begin{aligned}
\mathcal{M}_{b_1,\ldots,b_k}(a) &= \mathcal{E}(a|b_1,\ldots,b_k) = \int_{\Omega_a} a p(a|b_1,\ldots,b_k)\, da \\
\mathcal{V}_{b_1,\ldots,b_k}(a) &= \mathcal{E}\left(\|a - \mathcal{M}_{b_1,\ldots,b_k}(a)\|^2 | b_1,\ldots,b_k\right) \\
\mathcal{M}_{B_1,\ldots,B_k}(A) &= \mathcal{E}(A|B_1,\ldots,B_k) = \sum_{\Omega_A} A P(A|B_1,\ldots,B_k) \\
\mathcal{V}_{B_1,\ldots,B_k}(A) &= \mathcal{E}\left(\|A - \mathcal{M}(A)\|^2 | B_1,\ldots,B_k\right).
\end{aligned}$$

2) Modeling the Problem: Given a dictionary $\boldsymbol{D} \in \mathbb{R}^{n\times m}$, let Ω denote the set of all 2^m subdictionaries, where a subdictionary $\boldsymbol{S}$ will interchangeably be considered as a subset of the columns of $\boldsymbol{D}$ or as a matrix comprised of such columns. We assume that a random signal $\boldsymbol{x} \in \mathbb{R}^n$ is selected by the following process. With each subdictionary $\boldsymbol{S} \in \Omega$, we associate a

nonnegative probability $P(\boldsymbol{S})$, with $\sum_{\boldsymbol{S}\in\Omega} P(\boldsymbol{S}) = 1$. Furthermore, with each signal $\boldsymbol{x}$ in the range of $\boldsymbol{S}$ (that is, such that there exists a vector $\boldsymbol{z} \in \mathbb{R}^k$ satisfying $\boldsymbol{S}\boldsymbol{z} = \boldsymbol{x}$) denoted $\boldsymbol{x} \in \mathcal{R}(\boldsymbol{S})$, we associate a conditional pdf $p(\boldsymbol{x}|\boldsymbol{S})$. Then, the clean signal $\boldsymbol{x}$ is assumed to be generated by first randomly selecting $\boldsymbol{S}$ according to $P(\boldsymbol{S})$, and then randomly choosing $\boldsymbol{x} \in \boldsymbol{S}$ according to $p(\boldsymbol{x}|\boldsymbol{S})$. After the signal is generated, an additive random noise term $\boldsymbol{v}$, with pdf $p_v(\boldsymbol{v})$, is introduced, yielding a noisy signal $\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{v}$.

Note that $P(\boldsymbol{S})$ can be used to represent a tendency towards sparsity. For example, we can choose $P(\boldsymbol{S})$ to be a strongly decreasing function of the number of elements in $\boldsymbol{S}$, or we can choose $P(\boldsymbol{S})$ to be zero for all $\boldsymbol{S}$'s except those with a particular (small) number of elements.

Given $\boldsymbol{y}$, and assuming we know $p_v(\boldsymbol{v})$, $P(\boldsymbol{S})$, and $p(\boldsymbol{x}|\boldsymbol{S})$, our objective is to find an estimation $\hat{\boldsymbol{x}}$ that will be as close as possible to the clean signal $\boldsymbol{x}$ in some sense. In this work, we will mainly strive to minimize the conditional MSE

$$\mathrm{MSE}_y = \mathcal{E}\left(||\hat{\boldsymbol{x}} - \boldsymbol{x}||^2 \,|\boldsymbol{y}\right). \tag{4}$$

Note that typically one would expect to define the overall MSE without the condition over $\boldsymbol{y}$. However, this introduces a formidable yet unnecessary complication to the analysis that follows, and we shall avoid it.

3) Main Derivation: We first write the conditional MSE as the sum

$$\mathrm{MSE}_y = \sum_{\boldsymbol{S}\in\Omega} P(\boldsymbol{S}|\boldsymbol{y})\mathrm{MSE}_{S,y} \tag{5}$$

with $\mathrm{MSE}_{S,y}$ defined as

$$\mathrm{MSE}_{S,y} = \mathcal{E}\left(||\hat{\boldsymbol{x}} - \boldsymbol{x}||^2 \,|\boldsymbol{S},\boldsymbol{y}\right). \tag{6}$$

The term $\mathrm{MSE}_{S,y}$ is the MSE subject to a noisy signal $\boldsymbol{y}$ and a given subdictionary $\boldsymbol{S}$, and $P(\boldsymbol{S}|\boldsymbol{y})$ is the probability of $\boldsymbol{S}$ given a noisy signal $\boldsymbol{y}$. By Bayes's formula, the latter is given by

$$P(\boldsymbol{S}|\boldsymbol{y}) = \frac{p(\boldsymbol{y}|\boldsymbol{S})P(\boldsymbol{S})}{p(\boldsymbol{y})} \tag{7}$$

where

$$p(\boldsymbol{y}|\boldsymbol{S}) = \int_{\boldsymbol{x}\in\mathcal{R}(\boldsymbol{S})} p_v(\boldsymbol{y}-\boldsymbol{x})p(\boldsymbol{x}|\boldsymbol{S})\,d\boldsymbol{x} \tag{8}$$

is the pdf of $\boldsymbol{y}$ given the subdictionary $\boldsymbol{S}$.

Note that $p(\boldsymbol{y})$—the pdf of $\boldsymbol{y}$—can be computed directly or, more easily, obtained from the normalization requirement

$$\sum_{\boldsymbol{S}\in\Omega} P(\boldsymbol{S}|\boldsymbol{y}) = 1. \tag{9}$$

Nevertheless, as we shall soon see, it is not explicitly needed in our analysis.

Next, we consider $\mathrm{MSE}_{S,y}$, the MSE for a given $\boldsymbol{y}$ and subdictionary $\boldsymbol{S}$. Using the fact that $\mathcal{M}_{S,y}(\boldsymbol{x}) = \mathcal{E}(\boldsymbol{x}\,|\boldsymbol{S},\boldsymbol{y})$, we have

$$\begin{aligned}\mathcal{E}\left(||\boldsymbol{x}||^2\,|\boldsymbol{S},\boldsymbol{y}\right) &= \mathcal{E}\left(||\mathcal{M}_{S,y}(\boldsymbol{x}) + \boldsymbol{x} - \mathcal{M}_{S,y}(\boldsymbol{x})||^2\,|\boldsymbol{S},\boldsymbol{y}\right)\\ &= ||\mathcal{M}_{S,y}(\boldsymbol{x})||^2 + \mathcal{E}\left(||\boldsymbol{x} - \mathcal{M}_{S,y}(\boldsymbol{x})||^2\,|\boldsymbol{S},\boldsymbol{y}\right)\\ &= ||\mathcal{M}_{S,y}(\boldsymbol{x})||^2 + \mathcal{V}_{S,y}(\boldsymbol{x}).\end{aligned} \tag{10}$$

This well-known property, along with the linearity of the expectation, can be used to rewrite the first factor of the summation in (5) as follows:

$$\begin{aligned}\mathrm{MSE}_{S,y} &= \mathcal{E}\left(||\hat{\boldsymbol{x}} - \boldsymbol{x}||^2\,|\boldsymbol{S},\boldsymbol{y}\right)\\ &= \mathcal{E}\left(||\hat{\boldsymbol{x}}||^2 - 2\hat{\boldsymbol{x}}^T\boldsymbol{x} + ||\boldsymbol{x}||^2\,|\boldsymbol{S},\boldsymbol{y}\right)\\ &= ||\hat{\boldsymbol{x}}||^2 - 2\hat{\boldsymbol{x}}^T\mathcal{M}_{S,y}(\boldsymbol{x}) + ||\mathcal{M}_{S,y}(\boldsymbol{x})||^2 + \mathcal{V}_{S,y}(\boldsymbol{x})\\ &= ||\hat{\boldsymbol{x}} - \mathcal{M}_{S,y}(\boldsymbol{x})||^2 + \mathcal{V}_{S,y}(\boldsymbol{x}).\end{aligned} \tag{11}$$

Finally, plugging this into (5) we obtain

$$\begin{aligned}\mathrm{MSE}_y &= \sum_{\boldsymbol{S}\in\Omega}\left[||\hat{\boldsymbol{x}} - \mathcal{M}_{S,y}(\boldsymbol{x})||^2 + \mathcal{V}_{S,y}(\boldsymbol{x})\right]P(\boldsymbol{S}|\boldsymbol{y})\\ &= \mathcal{E}\left(||\hat{\boldsymbol{x}} - \mathcal{M}_{S,y}(\boldsymbol{x})||^2|\boldsymbol{y}\right) + \mathcal{E}\left(\mathcal{V}_{S,y}(\boldsymbol{x})|\boldsymbol{y}\right)\end{aligned} \tag{12}$$

with $P(\boldsymbol{S}|\boldsymbol{y})$ given by (7). As we have already mentioned, the overall MSE is given by

$$\mathrm{MSE} = \mathcal{E}\left(\mathrm{MSE}_y\right) = \int_{\boldsymbol{y}\in\mathbb{R}^n} \mathrm{MSE}_y p(\boldsymbol{y})\,d\boldsymbol{y} \tag{13}$$

but we shall not need this measure here.

4) The Optimal Estimator: By (12), the optimal $\hat{\boldsymbol{x}}$ that minimizes MSE_y is, not surprisingly, given by

$$\hat{\boldsymbol{x}}^{\mathrm{MMSE}} = \mathcal{E}\left(\mathcal{M}_{S,y}(\boldsymbol{x})|\boldsymbol{y}\right) \tag{14}$$

and, plugged to (12), the resulting optimal conditional MSE is given by

$$\begin{aligned}&\mathrm{MSE}_y^{\mathrm{MMSE}}\\ &= \mathcal{E}(||\mathcal{M}_{S,y}(\boldsymbol{x}) - \mathcal{E}(\mathcal{M}_{S,y}(\boldsymbol{x})|\boldsymbol{y})||^2|\boldsymbol{y}) + \mathcal{E}(\mathcal{V}_{S,y}(\boldsymbol{x})|\boldsymbol{y}).\end{aligned} \tag{15}$$

Finally, from (12) and (14) we obtain, for an arbitrary estimator $\hat{\boldsymbol{x}}$, the conditional MSE

$$\mathrm{MSE}_y = \mathrm{MSE}_y^{\mathrm{MMSE}} + ||\hat{\boldsymbol{x}} - \hat{\boldsymbol{x}}^{\mathrm{MMSE}}||^2. \tag{16}$$

This can be used to determine how much better the optimal estimator does compared to any other estimator.

5) The MAP Estimator: The MAP estimator is obtained by maximizing the probability of $\boldsymbol{x}$ given $\boldsymbol{y}$

$$\begin{aligned}\hat{\boldsymbol{x}}^{\mathrm{MAP}} &= \arg\max_{\boldsymbol{x}} p(\boldsymbol{x}|\boldsymbol{y})\\ &= \arg\max_{\boldsymbol{x}} \frac{p(\boldsymbol{y}|\boldsymbol{x})p(\boldsymbol{x})}{p(\boldsymbol{y})}\\ &= \arg\max_{\boldsymbol{x}} p_v(\boldsymbol{y}-\boldsymbol{x})p(\boldsymbol{x})\end{aligned} \tag{17}$$

with

$$p(\boldsymbol{x}) = \sum_{\boldsymbol{S}\in\Omega:\boldsymbol{x}\in\mathcal{R}(\boldsymbol{S})} p(\boldsymbol{x}|\boldsymbol{S})P(\boldsymbol{S}).$$

At the moment these expressions remain vague, but as we turn to use the specific signal and noise models discussed in Section III-A.2, these will assume an explicit form.

6) The Oracle: Suppose that the subdictionary $\boldsymbol{S}$ that was chosen in the generation of $\boldsymbol{x}$ is revealed to us. Given this information, we clearly minimize MSE_y by setting $\hat{\boldsymbol{x}} = \mathcal{M}_{S,y}(\boldsymbol{x})$

for the given $\boldsymbol{S}$. We call this the oracle estimator. The resulting conditional MSE is evidently given by the last term of (12)

$$\mathrm{MSE}_y^{\mathrm{oracle}} = \mathcal{E}\left(\mathcal{V}_{S,y}(\boldsymbol{x})|\boldsymbol{y}\right). \tag{18}$$

We shall use this estimator to assess the performance of the various alternatives and see how close we get to this "ideal" performance.

B. Back to Our Story—Sparse Representations

Our aim now is to harness the general derivation to the development of a practical algorithm for the sparse representation and white Gaussian noise. Motivated by the sparse-representation paradigm, we concentrate on the case where $P(\boldsymbol{S})$ depends only on the number of atoms (columns) in $\boldsymbol{S}$, denoted $|\boldsymbol{S}|$. We proceed with the basic case where $P(\boldsymbol{S})$ vanishes unless $|\boldsymbol{S}|$ is exactly equal to some particular $0 \le k \le \min(n,m)$, and $\boldsymbol{S}$ has column rank k. We denote the set of such $\boldsymbol{S}$'s by Ω_k, and define the uniform distribution

$$P(\boldsymbol{S}) = \begin{cases} \frac{1}{|\Omega_k|}, & \boldsymbol{S} \in \Omega_k \\ 0, & \text{otherwise.} \end{cases}$$

We assume throughout that the columns of $\boldsymbol{D}$ are normalized, $\|\boldsymbol{d}_j\| = 1$, for $j = 1, \ldots, m$. This assumption serves only to simplify the expressions we are about to obtain. Next, we recall that the noise is modeled via a Gaussian distribution with zero mean and variance σ^2, and thus

$$p\left(\boldsymbol{y}|\boldsymbol{x}\right) = p_v\left(\boldsymbol{y}-\boldsymbol{x}\right) = \frac{1}{\left(2\pi\sigma^2\right)^{n/2}} \cdot \exp\left\{\frac{-\|\boldsymbol{y}-\boldsymbol{x}\|^2}{2\sigma^2}\right\}. \tag{19}$$

Similarly, given the subdictionary $\boldsymbol{S}$ from which $\boldsymbol{x}$ is drawn, the signal $\boldsymbol{x}$ is assumed to be generated via a Gaussian distribution with mean zero and variance σ_x^2, thus $p(\boldsymbol{x}|\boldsymbol{S})$ is given by

$$p\left(\boldsymbol{x}|\boldsymbol{S}\right) = \begin{cases} \frac{1}{\left(2\pi\sigma_x^2\right)^{k/2}} \cdot \exp\left\{\frac{-\|\boldsymbol{x}\|^2}{2\sigma_x^2}\right\}, & \boldsymbol{x} \in \mathcal{R}(\boldsymbol{S}) \\ 0, & \text{otherwise.} \end{cases} \tag{20}$$

Note that this distribution does not align with the intuitive creation of $\boldsymbol{x}$ as $\boldsymbol{S}\boldsymbol{z}$ with a Gaussian vector $\boldsymbol{z} \in \mathbb{R}^k$ with i.i.d. entries. Instead, we assume that an orthogonalized basis for this subdictionary has been created and then multiplied by $\boldsymbol{z}$. We adopt the latter model for simplicity; the former model has also been worked out in full, but we omit it here because it is somewhat more complicated and seems to afford only modest additional insights.

For convenience, we introduce the notation $c^2 = \sigma_x^2/(\sigma^2 + \sigma_x^2)$ (cf. (2)). Also, we denote the orthogonal projection of any vector $\boldsymbol{a}$ onto the subspace spanned by the columns of $\boldsymbol{S}$ by

$$\boldsymbol{a}_S = \boldsymbol{S}\left(\boldsymbol{S}^T\boldsymbol{S}\right)^{-1}\boldsymbol{S}^T\boldsymbol{a}.$$

We now follow the general derivation given above. From (8), we can develop a closed-form expression for $p(\boldsymbol{y}|\boldsymbol{S})$. By integration and rearrangement we obtain (21), shown at the bottom of the page. Since the only dependence of $p(\boldsymbol{y}|\boldsymbol{S})$ on $\boldsymbol{S}$ is through the rightmost factor, we immediately obtain by (7) and (9) the simple formula

$$P(\boldsymbol{S}|\boldsymbol{y}) = \frac{\exp\left\{-\frac{c^2\|\boldsymbol{y}-\boldsymbol{y}_S\|^2}{2\sigma^2}\right\}}{\sum_{\boldsymbol{S}'\in\Omega_k}\exp\left\{\frac{c^2\|\boldsymbol{y}-\boldsymbol{y}_{S'}\|^2}{2\sigma^2}\right\}}. \tag{22}$$

The denominator here is just a normalization. The numerator implies that, given a noisy signal $\boldsymbol{y}$, the probability that the clean signal was selected from the subspace $\boldsymbol{S}$ decays at a Gaussian rate with the distance between $\boldsymbol{y}$ and $\boldsymbol{S}$, i.e., $\|\boldsymbol{y}-\boldsymbol{y}_S\|$. This result is expected, given the Gaussian noise distribution.

Continuing to follow the general analysis, we compute the conditional mean, $\mathcal{M}_{S,y}(\boldsymbol{x})$, for which we require the conditional probability

$$\begin{aligned} p(\boldsymbol{x}|\boldsymbol{S},\boldsymbol{y}) &= \frac{p(\boldsymbol{y}|\boldsymbol{S},\boldsymbol{x})\,p(\boldsymbol{x}|\boldsymbol{S})}{p(\boldsymbol{y}|\boldsymbol{S})} \\ &= \frac{1}{p(\boldsymbol{y}|\boldsymbol{S})}\cdot\frac{1}{(2\pi\sigma^2)^{n/2}}\cdot\exp\left\{\frac{-\|\boldsymbol{y}-\boldsymbol{x}\|^2}{2\sigma^2}\right\} \\ &\quad\cdot\frac{1}{(2\pi\sigma_x^2)^{k/2}}\cdot\exp\left\{\frac{-\|\boldsymbol{x}\|^2}{2\sigma_x^2}\right\}. \end{aligned} \tag{23}$$

By integration, we then obtain the simple result

$$\mathcal{M}_{S,y}(\boldsymbol{x}) = \int_{\boldsymbol{x}\in\mathcal{R}(\boldsymbol{S})}\boldsymbol{x}p(\boldsymbol{x}|\boldsymbol{S},\boldsymbol{y})d\boldsymbol{x} = c^2\boldsymbol{y}_S. \tag{24}$$

Now the conditional variance can be computed, yielding

$$\mathcal{V}_{S,y}(\boldsymbol{x}) = \int_{\boldsymbol{x}\in\mathcal{R}(\boldsymbol{S})}\|\boldsymbol{x}-c^2\boldsymbol{y}_S\|^2 p(\boldsymbol{x}|\boldsymbol{S},\boldsymbol{y})d\boldsymbol{x} = kc^2\sigma^2 \tag{25}$$

which is independent of $\boldsymbol{S}$ and $\boldsymbol{y}$. Thus, the oracle MSE_y in this case is simply

$$\mathrm{MSE}_y^{\mathrm{oracle}} = kc^2\sigma^2. \tag{26}$$

$$\begin{aligned} p(\boldsymbol{y}|\boldsymbol{S}) &= \int_{\boldsymbol{x}\in\mathcal{R}(\boldsymbol{S})}\frac{1}{\left(2\pi\sigma^2\right)^{n/2}\cdot\left(2\pi\sigma_x^2\right)^{k/2}}\cdot\exp\left\{\frac{-\|\boldsymbol{y}-\boldsymbol{x}\|^2}{2\sigma^2}+\frac{-\|\boldsymbol{x}\|^2}{2\sigma_x^2}\right\}d\boldsymbol{x} \\ &= \frac{(1-c^2)^{k/2}}{|\Omega_k|\left(2\pi\sigma^2\right)^{n/2}}\cdot\exp\left\{\frac{-(1-c^2)\|\boldsymbol{y}\|^2}{2\sigma^2}\right\}\cdot\exp\left\{\frac{-c^2\|\boldsymbol{y}-\boldsymbol{y}_S\|^2}{2\sigma^2}\right\}. \end{aligned} \tag{21}$$

The optimal estimator is given, by (14), as

$$\begin{aligned}\hat{\boldsymbol{x}}^{\mathrm{MMSE}} &= c^2 \sum_{\boldsymbol{S}\in\Omega_k} \boldsymbol{y}_S P(\boldsymbol{S}|\boldsymbol{y}) \\ &= \frac{c^2}{\sum_{\boldsymbol{S}'\in\Omega_k} \exp\left\{-\frac{c^2\|\boldsymbol{y}-\boldsymbol{y}_{S'}\|^2}{2\sigma^2}\right\}} \\ &\quad \cdot \sum_{\boldsymbol{S}\in\Omega_k} \exp\left\{\frac{-c^2\|\boldsymbol{y}-\boldsymbol{y}_S\|^2}{2\sigma^2}\right\} \boldsymbol{y}_S \end{aligned} \tag{27}$$

with $P(\boldsymbol{S}|\boldsymbol{y})$ taken from (22). This MMSE estimate is a weighted average of the projections of $\boldsymbol{y}$ onto all the possible subspaces $\boldsymbol{S} \in \Omega_k$, as claimed. The MSE of this estimate is given by

$$\mathrm{MSE}_y^{\mathrm{MMSE}} = kc^2\sigma^2 + \sum_{\boldsymbol{S}\in\Omega_k} \|\hat{\boldsymbol{x}}^{\mathrm{MMSE}} - c^2\boldsymbol{y}_S\|^2 P(\boldsymbol{S}|\boldsymbol{y}). \tag{28}$$

The latter can also be written as

$$\mathrm{MSE}_y^{\mathrm{MMSE}} = kc^2\sigma^2 - \|\hat{\boldsymbol{x}}^{\mathrm{MMSE}}\|^2 + \sum_{\boldsymbol{S}\in\Omega_k} \|c^2\boldsymbol{y}_S\|^2 P(\boldsymbol{S}|\boldsymbol{y}). \tag{29}$$

We remark that *any* spherically symmetric $p_v(\boldsymbol{v})$ and $p(\boldsymbol{x}|\boldsymbol{S})$ produce a conditional mean, $\mathcal{M}_{S,y}(\boldsymbol{x})$, that is equal to $\boldsymbol{y}_S$ times some scalar coefficient. The choice of Gaussian distributions makes the result in (24) particularly simple in that the coefficient c^2 is independent of $\boldsymbol{y}$ and $\boldsymbol{S}$.

Next, we consider the MAP estimator, using (17). For simplicity, we shall neglect the fact that some $\boldsymbol{x}$'s may lie on intersections of two or more subdictionaries in Ω_k, and therefore their pdf is higher according to our model. This is a set of measure zero, and it therefore does not influence the MMSE solution, but it does influence somewhat the MAP solution for $\boldsymbol{y}$'s that are close to such $\boldsymbol{x}$'s. We can overcome this technical difficulty by modifying our model slightly so as to eliminate the favoring of such $\boldsymbol{x}$'s. Noting that $P(\boldsymbol{S})$ is a constant for all $\boldsymbol{S} \in \Omega_k$, we obtain from (17)

$$\hat{\boldsymbol{x}}^{\mathrm{MAP}} = \arg\max_{\boldsymbol{x}\in\mathcal{R}(\Omega_k)} \exp\left\{\frac{-\|\boldsymbol{y}-\boldsymbol{x}\|^2}{2\sigma^2}\right\} \cdot \exp\left\{\frac{-\|\boldsymbol{x}\|^2}{2\sigma_x^2}\right\} \tag{30}$$

where $\mathcal{R}(\Omega_k)$ is defined as the union of the ranges of all $\boldsymbol{S} \in \Omega_k$. Multiplying through by $\exp(2c^2\sigma^2)$, we find that the maximum is obtained by minimizing $c^2\|\boldsymbol{y}-\boldsymbol{x}\|^2 + (1-c^2)\|\boldsymbol{x}\|^2$, subject to the constraint that $\boldsymbol{x}$ belongs to some $\boldsymbol{S} \in \Omega_k$. The resulting estimator is readily found to be given by

$$\hat{\boldsymbol{x}}^{\mathrm{MAP}} = c^2\boldsymbol{y}_{S_{\mathrm{MAP}}}, \tag{31}$$

where $\boldsymbol{S}_{\mathrm{MAP}}$ is the subspace $\boldsymbol{S} \in \Omega_k$ which is closest to $\boldsymbol{y}$, i.e., for which $\|\boldsymbol{y}-\boldsymbol{y}_S\|^2$ is the smallest. The resulting MSE_y is given by substituting $\hat{\boldsymbol{x}}^{\mathrm{MAP}}$ for $\hat{\boldsymbol{x}}$ in (16).

Note that in all the estimators we derive, the oracle, the MMSE, and the MAP, there is a factor of c^2 that performs a shrinking of the estimate. For the model of $\boldsymbol{x}$ chosen, this is a mandatory step that was omitted in Section II.

C. Combining it All

It is now time to combine the theoretical analysis of the section and the estimators we tested in Section II. We have several goals in this discussion.

- We would like to evaluate both the expressions and the empirical values of the MSE for the oracle, the MMSE, and the MAP estimators, and show their behavior as a function of the input noise power σ.
- We would like to show how the above aligns with the actual OMP and the RandOMP results obtained.
- This discussion will help explain two choices made in the RandOMP algorithm—the rule for drawing the next atom, and the requirement of a plain averaging of the representations.

We start by constructing a synthetic experiment that enables a comparison of the various estimators and formulas developed. We build a random dictionary of size 20×30 with ℓ_2-normalized columns. We generate signals following the model described above by randomly choosing a support with k columns (we vary in the range $[1,3]$), orthogonalizing the chosen columns, and multiplying them by a random i.i.d. vector with entries drawn from $N(0,1)$ (i.e., $\sigma_x = 1$). We add noise to these signals with σ in the range $[0.1, 2]$ and evaluate the following values.

1. **Empirical Oracle** estimation and its MSE. This estimator is simply the projection of $\boldsymbol{y}$ on the correct support, followed by a multiplication by c^2, as described in (24).
2. **Theoretical Oracle** estimation error, as given in (26).
3. **Empirical MMSE** estimation and its MSE. We use the formula in (27) in order to compute the estimation, and then assess its error empirically. Note that in applying this formula we gather all the $\binom{30}{k}$ possible supports, compute the projection of $\boldsymbol{y}$ onto them, and weight them according to the formula. This explains why in the experiment reported here we have restricted the sizes involved.
4. **Theoretical MMSE** estimation error, using (29) directly.
5. **Empirical MAP** estimation and its MSE. We use the analytic solution to (30) as described above, by sweeping through all the possible supports, and searching the one with the smallest projection error. This gives us the MAP estimation, and its error is evaluated empirically.
6. **Theoretical MAP** estimation error, as given in (16), when plugging in the MAP estimation.
7. **OMP** estimation and its MSE. The OMP is the same as described in Section II, but the stopping rule is based on the knowledge of k, rather than on representation error. Following the MAP analysis done in Section III, the result is multiplied by c^2 as well.
8. **Averaged RandOMP** estimation and its MSE. The algorithm generates $J_0 = 100$ representations and averages them. As in the OMP, the stopping rule for those is the number of atoms k, and the result is also multiplied by c^2.

The above process is averaged over 1000 signal generations, and the resulting noise attenuation values are shown in Fig. 7 for $k = 1,\ 2,$ and 3. As in earlier figures, we present the ratio between the actual MSE obtained and the input noise power.

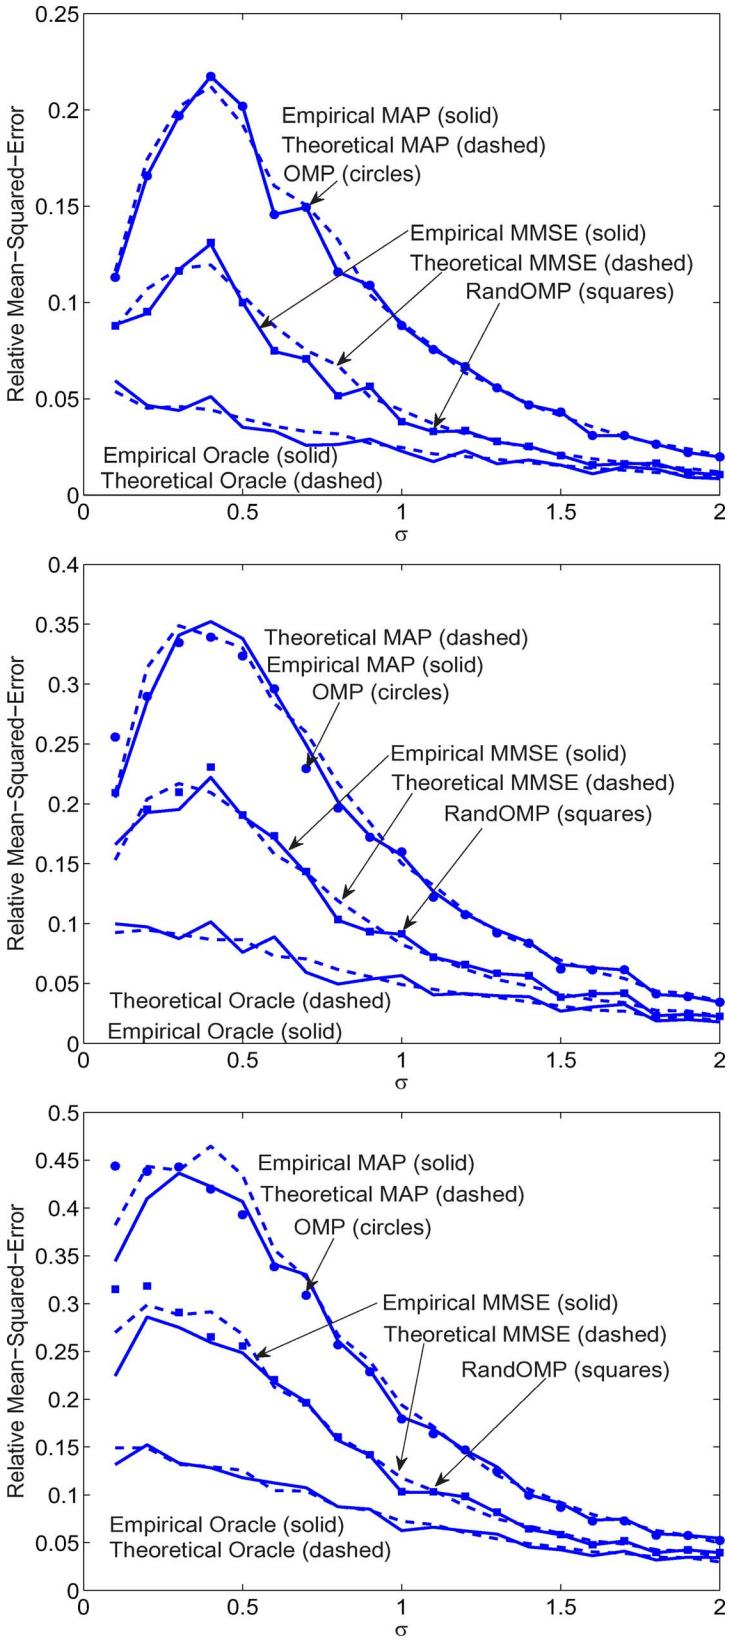

Fig. 7. MSE as a function of the input noise for $k = 1$ (top), $k = 2$ (middle), and $k = 3$ (bottom).

First we draw attention to several general observations. As expected, we see in all these graphs that there is a good alignment between the theoretical and the empirical evaluation of the MSE for the oracle, the MMSE, and the MAP estimators. In fact, since the analysis is exact for this experiment, the differences are only due to the finite number of tests per σ. We also see that the de-

noising performance weakens as k grows. A third and intriguing observation that we will not explore here is the fact that there appears to be a critical input noise power ($\sigma \approx 0.4$) for which the MAP and the MMSE estimators (and their approximations) give their worst denoising performance, as exhibited by the hump in all the MMSE/MAP cases.

The OMP algorithm is an attempt to approximate the MAP estimation, replacing the need for sweeping through all the possible supports by a greedy detection of the involved atoms. As such, we expect it to be competitive and close to the MAP results we get (either analytically or empirically). In fact, for $k = 1$, it aligns perfectly with the empirical MAP, since both are going through the same computational stages. As k grows, there are some differences between the empirical MAP and the OMP, and especially for low noise, but for the cases studied here these differences are relatively small.

Just as OMP is an attempt to approximate the MAP estimation, the RandOMP averaging is approximating the MMSE estimator, thereby yielding much better denoising than OMP. The core idea is to replace the summation over all possible supports with a much smaller selected group of representations that are sampled from the distribution $P(\boldsymbol{S}|\boldsymbol{y})$ governing by the weights in (27). Indeed, the representations chosen by RandOMP are those that correspond to large weights, since they are built in a way that leads to small projection error $\|\boldsymbol{y} - \boldsymbol{y}_P\|^2$ for the k atoms chosen. Since the sampling already mimics approximately the required distribution, all that remains is a simple averaging, as indeed we do in practice. What is required is to tune the sampling to be faithful, and for that we revisit the case of $k = 1$.

Considering the case of $k = 1$, we see from (27) that an atom should be chosen as a candidate representation with a probability proportional to $\exp\{-c^2\|\boldsymbol{y} - \boldsymbol{y}_P\|^2/2\sigma^2\}$. This in turn implies that this probability is also proportional to[4] $\exp\{c^2|\boldsymbol{y}^T\boldsymbol{d}_i|^2/2\sigma^2\}$. Thus, RandOMP as described in Section II is in perfect agreement with this probability if the representation's cardinality is $k = 1$, and this explains the goodness of fit of RandOMP with the empirical MSE in Fig. 7—top. However, we also see that RandOMP remains close to the empirical MMSE for $k = 2$ and 3, implying that while our sampling strategy is not perfect for $k > 1$, it is fair enough. Further investigation is required to better sample the representations in order to get closer to the MSE estimate.

We note an additional advantage of RandOMP: the MMSE estimator varies continuously with y, whereas the MAP estimator does not, possibly leading to artifacts.

All the above discussion is heavily based on the Gaussianity assumptions made on both the representation's coefficients and the noise. These assumptions enabled getting relatively simple closed-form expressions in our analysis, and in the formula with which RandOMP draws its atoms. Nevertheless, much of this work can be extended to non-Gaussian distributions, by following the road map laid out in this paper, though it might be technically more difficult. Along with such modifications in the analysis, the RandOMP algorithm should change as well, as a direct consequence of the expected change in $P(\boldsymbol{S}|\boldsymbol{y})$, and this is in principle computable. Alternatively, one could proceed using the Random-OMP with the Gaussian formula, in the belief that there is some robustness in the algorithm to changes in the noise model. Preliminary tests along these lines provide promising results, but we leave these matters to a separate work.

As a last point in this section, we return to our observation in Section II-C regarding the nonsparse representation that the MMSE approximation provides. The experiments done in this section also lead to dense representations, with either the exact MMSE or its RandOMP approximation. It is clear that there is an excellent sparse solution to our estimation task—this is the oracle. However, this estimator is unattainable. The MAP (or better yet, the OMP which is its practical approximation) provides a sparse solution, but its quality is inferior in ℓ_2-norm, compared to the MMSE. We also know that the MSE estimation is the best we can hope for.

In some applications, alongside with the desired accuracy of the estimation, one desires a sparse outcome. This, for example, is relevant for compression or recognition purposes. The question is: Could there be a practically attainable sparse solution that leads to near-optimal MSE? Clearly, we cannot hope to bypass the MMSE in terms of the error, but could we get close to it, bypassing the performance of the MAP? The answer to this question is positive. Once we have estimated the MMSE using RandOMP, we can then project it to a sparse solution by running a regular OMP on it and demanding a representation with k non-zeros. Surprisingly enough, this does not lead to MAP, but rather to a very good (and sparse!) approximation of the MMSE. Fig. 8 presents such an experiment, which is a direct continuation of the one reported in Fig. 7. For the case of $k = 3$, we test the OMP, RandOMP, exact MMSE (exhaustively computed), and the projected version of RandOMP to three atoms, and average over 1000 experiments. As can be seen, the performance of this sparse solution is far better than the MAP, and very close to the MMSE one.

D. Summary

Under the assumptions of this section, we obtain simple explicit expressions for the optimal (MMSE) estimator and its resulting MSE$_y$. The optimal estimator turns out to be a weighted average of the orthogonal projections of the noisy signal on the feasible subspaces, multiplied by a "shrinkage factor" c^2, which tends to zero when the noise variance σ^2 is large compared to the signal variance, σ_x^2, and to 1 when the opposite is true. The weights in the weighted average depend on the distances between $\boldsymbol{y}$ and the subspaces, favoring short distances of course, especially when $c^2\|\boldsymbol{y}\|^2/\sigma^2$ is large.

While the expressions obtained are indeed simple, they involve either an intolerable summations over $\binom{m}{k}$ (for the MMSE estimate), or searching over this amount of subspaces (for the MAP). Thus, these formulas are impractical for direct use. In that sense, one should consider the RandOMP approach in Section II as a sampler from this huge set of subspaces over which we average. Roughly speaking, since the RandOMP algorithm tends to find near—by subspaces that lead to sparse representations, it gives priority to elements in the summation in (27) that are assigned higher weights. We see experimentally that RandOMP samples well from the representations, judging by the

[4]Since the columns of the dictionary are normalized, the projection is given by $\boldsymbol{y}_P = (\boldsymbol{y}^T\boldsymbol{d}_i)\cdot\boldsymbol{d}_i$. Thus, $\|\boldsymbol{y} - \boldsymbol{y}_P\|^2 = \|\boldsymbol{y}\|^2 - (\boldsymbol{y}^T\boldsymbol{d}_i)^2$. The term $\exp\{-c^2\|\boldsymbol{y}\|^2\}$ is therefore a constant that cancels out in the normalization.

proximity of its results to the MMSE error (both empirical and theoretical).

The results of this section can easily be extended to the case where we allow a range of values of k with given probabilities. That is, we can extend these results for the case where $P(\boldsymbol{S}) = f(|\boldsymbol{S}|)$, for general nonnegative functions f.

IV. Summary and Conclusion

The Orthogonal Matching Pursuit (OMP) is a simple and fast algorithm for approximating the sparse representation for a given signal. It can be used for denoising of signals, as a way to approximate the MAP estimation. In this work, we have shown that by running this algorithm several times in a slightly modified version that randomizes its outcome, one can obtain a collection of competing representations, and those can be averaged to lead to far better denoising performance. This work starts by showing how to obtain a set of such representations to merge, how to combine them wisely, and what kind of results to expect. The analytic part of this paper explains this averaging as a way to approximate the MMSE estimate as a sampler of the summation required. Future work on this topic should consider better sampling strategies for better approximation of the MMSE result, an analytical and numerical study of the required number of samples, an assessment of the robustness of this approach with respect to non-Gaussian distribution of signals and limited accuracy in determining their variance, and exploration of special cases for which practical deterministic algorithms are within reach.

References

[1] S. Baker and T. Kanade, "Limits on super-resolution and how to break them," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 24, no. 9, pp. 1167–1183, Sep. 2002.

[2] A. M. Bruckstein, D. L. Donoho, and M. Elad, "From sparse solutions of systems of equations to sparse modeling of signals and images," *SIAM Rev.*, vol. 51, no. 1, pp. 34–81, Feb. 2009.

[3] A. Buades, B. Coll, and J. M. Morel, "A review of image denoising algorithms, with a new one," *Multiscale Modeling and Simulation*, vol. 4, no. 2, pp. 490–530, 2005.

[4] S. Chen, S. A. Billings, and W. Luo, "Orthogonal least squares methods and their application to non-linear system identification," *Int. J. Contr.*, vol. 50, no. 5, pp. 1873–96, 1989.

[5] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM J. Scient. Comput.*, vol. 20, no. 1, pp. 33–61, 1998.

[6] G. Davis, S. Mallat, and M. Avellaneda, "Adaptive greedy approximations," *J. Constructive Approx.*, vol. 13, pp. 57–98, 1997.

[7] G. Davis, S. Mallat, and Z. Zhang, "Adaptive time-frequency decompositions," *Opt. Eng.*, vol. 33, no. 7, pp. 2183–91, 1994.

[8] D. Datsenko and M. Elad, "Example-based single image super-resolution: A global MAP approach with outlier rejection," *J. Multidim. Syst. Signal Process.*, vol. 18, no. 2, pp. 103–121, Sep. 2007.

[9] D. L. Donoho, "For most large underdetermined systems of linear equations, the minimal ℓ_1-norm near-solution approximates the sparsest near-solution," *Commun. Pure Appl. Math.*, vol. 59, no. 7, pp. 907–934, Jul. 2006.

[10] D. L. Donoho and M. Elad, "On the stability of the basis pursuit in the presence of noise," *Signal Process.*, vol. 86, no. 3, pp. 511–532, Mar. 2006.

[11] D. L. Donoho, M. Elad, and V. Temlyakov, "Stable recovery of sparse overcomplete representations in the presence of noise," *IEEE Trans. Inf. Theory*, vol. 52, no. 1, pp. 6–18, Jan. 2006.

[12] M. Elad and M. Aharon, "Image denoising via sparse and redundant representations over learned dictionaries," *IEEE Trans. Image Process.*, vol. 15, no. 12, pp. 3736–3745, Dec. 2006.

[13] A. K. Fletcher, S. Rangan, V. K. Goyal, and K. Ramchandran, "Analysis of denoising by sparse approximation with random frame asymptotics," in *Proc. IEEE Int. Symp. Information Theory*, Adelaide, Australia, Sep. 2005, pp. 1706–1710.

[14] A. K. Fletcher, S. Rangan, V. K. Goyal, and K. Ramchandran, "Denoising by sparse approximation: Error bounds based on rate-distortion theory," *EURASIP J. Appl. Signal Process.*, 2006, paper no. 26318.

[15] W. T. Freeman, T. R. Jones, and E. C. Pasztor, "Example-based super-resolution," *IEEE Comput. Graphics and Applic.*, vol. 22, no. 2, pp. 56–65, 2002.

[16] W. T. Freeman, E. C. Pasztor, and O. T. Carmichael, "Learning low-level vision," *Int. J. Comput. Vision*, vol. 40, no. 1, pp. 25–47, 2000.

[17] J. J. Fuchs, "Recovery of exact sparse representations in the presence of bounded noise," *IEEE Trans. Inf. Theory*, vol. 51, no. 10, pp. 3601–3608, Oct. 2005.

[18] R. Gribonval, R. Figueras, and P. Vandergheynst, "A simple test to check the optimality of a sparse signal approximation," *Signal Process.*, vol. 86, no. 3, pp. 496–510, Mar. 2006.

[19] E. G. Larsson and Y. Selen, "Linear gregression with a sparse parameter vector," *IEEE Trans. Signal Process.*, vol. 55, no. 2, pp. 451–460, Feb. 2007.

[20] S. Mallat, *A Wavelet Tour of Signal Processing*, 3rd ed. San Diego, CA: Academic, 2009.

[21] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. Signal Process.*, vol. 41, no. 12, pp. 3397–3415, Dec. 1993.

[22] B. K. Natarajan, "Sparse approximate solutions to linear systems," *SIAM J. Comput.*, vol. 24, pp. 227–234, 1995.

[23] Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Proc. 27th Asilomar Conf. Signals, Systems and Computers*, Thousand Oaks, CA, 1993, vol. 1, pp. 40–44.

[24] L. C. Pickup, S. J. Roberts, and A. Zisserman, "A sampled texture prior for image super-resolution," *Adv. Neural Inf. Process. Syst.*, vol. 16, pp. 1587–1594, 2003.

[25] P. Schnitter, L. C. Potter, and J. Ziniel, "Fast Bayesian matching pursuit," in *Proc. Workshop on Information Theory and Applications (ITA)*, La Jolla, CA, Jan. 2008.

[26] P. Schniter, L. C. Potter, and J. Ziniel, "Fast Bayesian matching pursuit: Model uncertainty and parameter estimation for sparse linear models," *IEEE Trans. Signal Process.*, submitted for publication.

[27] J. A. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Trans. Inf. Theory*, vol. 50, no. 10, pp. 2231–2242, Oct. 2004.

[28] J. A. Tropp, "Just relax: Convex programming methods for subset selection and sparse approximation," *IEEE Trans. Inf. Theory*, vol. 52, no. 3, pp. 1030–1051, Mar. 2006.

[29] B. Wohlberg, "Noise sensitivity of sparse signal representations: Reconstruction error bounds for the inverse problem," *IEEE Trans. Signal Process.*, vol. 51, no. 12, pp. 3053–3060, Dec. 2003.

Michael Elad (M'98–SM'08) received the B.Sc degree in 1986, the M.Sc. degree (supervised by Prof. David Malah) in 1988, and the D.Sc. degree (supervised by Prof. Arie Feuer) in 1997, all from the Department of Electrical Engineering at the Technion–Israel Institute of Technology, Haifa, Israel.

From 1988 to 1993, he served in the Israeli Air Force. From 1997 to 2000, he worked at Hewlett-Packard Laboratories as an R&D engineer. From 2000 to 2001, he headed the research division at Jigami Corporation, **(city?)** Israel. During the years 2001–2003, he spent a postdoctoral period as a Research Associate with the Computer Science Department at Stanford University (SCCM program), Stanford, CA. In September 2003, he returned to the Technion, assuming a tenure-track assistant professorship position in the Department of Computer Science. In May 2007, he was tenured to an associate professorship. He works in the field of signal and image processing, specializing in particular on inverse problems, sparse representations, and overcomplete transforms.

Prof. Elad received the Technion's Best Lecturer award six times (1999, 2000, 2004, 2005, 2006, and 2009), he is the recipient of the Solomon Simon Mani award for excellence in teaching in 2007, and he is also the recipient of the Henri Taub Prize for academic excellence (2008). He is currently serving as an Associate Editor for both IEEE Transactions on Image Processing and *SIAM Journal on Imaging Sciences* (SIIMS).

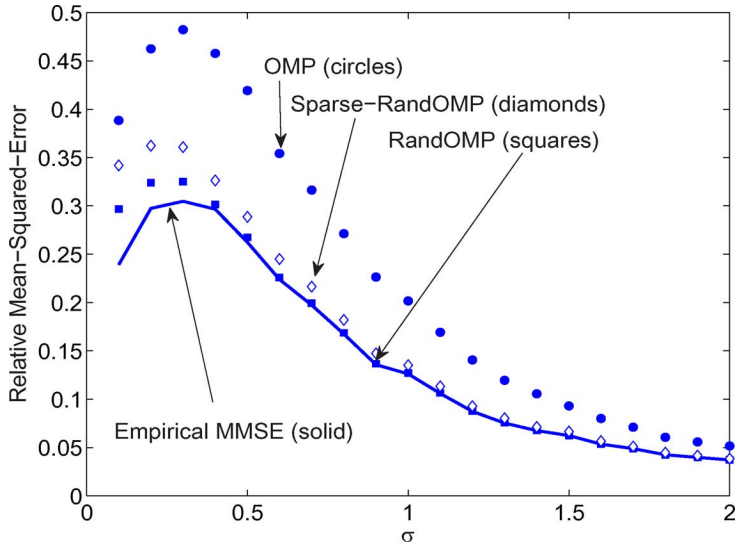

Fig. 8. MSE as a function of the input noise for a sparse solution obtained by projecting the RandOMP solution to $k = 3$ atoms.

Irad Yavneh received the Ph.D. degree in applied mathematics from the Weizmann Institute of Science, Rehovot, Israel, in 1991.

He is a Professor at the Faculty of Computer Science, Technion–Israel Institute of Technology, Haifa, Israel. His research interests include multiscale computational techniques, scientific computing and computational physics, image processing and analysis, and geophysical fluid dynamics.