

Leopold-Franzens-Universität Innsbruck

Preprint-Series: Department of Mathematics - Applied Mathematics

# Sequential Learning of Analysis Operators

Michael Sandbichler and Karin Schnass



Technikerstraße 13 - 6020 Innsbruck - Austria Tel.: +43 512 507 53803 Fax: +43 512 507 53898 https://applied-math.uibk.ac.at

Nr. 35 20. April 2017

# Sequential Learning of Analysis Operators

Michael Sandbichler & Karin Schnass

Abstract—In this paper two sequential algorithms for learning analysis operators are presented. They are built upon the same optimisation principle underlying both Analysis K-SVD and Analysis SimCO and use a stochastic gradient descent approach similar to ASimCO. The sequential analysis operator learning (SAOL) algorithm is based on projected gradient descent with an appropriately chosen step size while the implicit SAOL (ISAOL) algorithm avoids choosing a step size altogether by using a strategy inspired by the implicit Euler scheme for solving ordinary differential equations. Both algorithms are tested on synthetic and image data in comparison to Analysis SimCO and found to give slightly better recovery rates resp. decay of the objective function. In a final denoising experiment the presented algorithms are again shown to perform well in comparison to the state of the art algorithm ASimCO.

*Index Terms*—analysis operator learning, analysis dictionary learning, online learning, cosparse, sequential, stochastic gradient descent, thresholding, denoising

#### I. INTRODUCTION

Many tasks in high dimensional signal processing, such as denoising or reconstruction from incomplete information, can be efficiently solved if the data at hand is known to have intrinsic low dimension. One popular model with intrinsic low dimension is the union of subspaces model, where every signal is assumed to lie in one of the low dimensional linear subspaces. However, as the number of subspaces increases, the model becomes more and more cumbersome to use unless the subspaces can be parametrised. Two examples of large unions of parametrised subspaces, that have been successfully employed, are sparsity in a dictionary and cosparsity in an analysis operator. In the sparse model the subspaces correspond to the linear span of just a few normalised columns, also known as atoms, from a  $K \times d$  dictionary matrix,  $\Phi = (\phi_1 \dots \phi_K)$  with  $\|\phi_k\|_2 = 1$ , meaning, any data point y can be approximately represented as superposition of  $S \ll d$  dictionary elements. If we denote the restriction of the dictionary to the atoms/columns indexed by I as  $\Phi_I$ , we have

$$y \in \bigcup_{|I| \le S} \operatorname{colspan} \Phi_I$$
, or  $y \approx \Phi x$ , with x sparse.

In the cosparse model the subspaces correspond to the orthogonal complement of the span of some normalised rows, also known as analysers, from a  $d \times K$  analysis operator  $\Omega = (\omega_1^* \dots \omega_K^*)^*$  with  $\|\omega_k\|_2 = 1$ . This means that any data point y is orthogonal to  $\ell$  analysers or in other words that the vector  $\Omega y$  has  $\ell$  zero entries and is sparse. If we denote the restriction of the analysis operator to the analysers/rows indexed by J as  $\Omega_J$ , we have

$$y \in \bigcup_{|J| \ge \ell} (\operatorname{rowspan} \Omega_J)^{\perp}$$
, or  $\Omega y \approx z$ , with z sparse.

However, before being able to exploit these models for a given data class, it is necessary to identify the parametrising dictionary or analysis operator. This can be done either via a theoretical analysis or a learning approach. While dictionary learning is by now an established field, see [13] for an introductory survey, results in analysis operator learning are still countable, [18], [12], [11], [19], [10], [14], [2], [4], [7], [17], [3].

**Contribution:** In this work we will contribute to the development of the field by developing two sequential algorithms for learning analysis operators, which outperform state of the art algorithms such as Analysis K-SVD, [14] and Analysis SimCo, [2], [3], in terms of convergence speed while retaining the same performance.

Outline: The paper is organised as follows. After introducing the necessary notation, in the next section we will remotivate the opimisation principle that is the starting point of A-KSVD and ASimCO and shortly discuss the advantages and disadvantages of the two algorithms. We will then take a gradient descent approach similar to ASimCO, replacing the costly line search with a clever choice for the step size and a projection to stay sequential, and test our algorithm both on synthetic and image data. Inspired by the Euler scheme for solving ordinary differential equations, in Section III we will invest a little in the memory requirements of our algorithm in return for avoiding the stepsize alltogether. After testing the algorithm on synthetic and image data and comparing both our algorithms to ASimCO, in Section V we apply them to image denoising again in comparison to ASimCO. Finally, in the last section we provide a short discussion of our results and point out future directions of research. Notation: Before finally hitting the slopes, we summarise the notational conventions used throughout this paper. The operators  $\Omega$  and  $\Gamma$  will always denote matrices in  $\mathbb{R}^{K \times d}$ and for a matrix A we denote its transpose by  $A^*$ . More specifically, we will mostly consider matrices in the oblique manifold  $\mathcal{A} := \{ \Gamma \in \mathbb{R}^{K \times d} : \forall k \in [K] : \|\gamma_k\|_2 = 1 \}$ , where  $\gamma_k$  denotes the k-th row of the matrix  $\Gamma$ . By [n], we denote the set  $\{1, 2, \ldots, n\}$  and we adopt the standard notation |M|for the cardinality of a set M. By  $\Gamma_J$  with  $J \subset [K]$  we denote the restriction of  $\Gamma$  to the rows indexed by J.

A vector  $y \in \mathbb{R}^d$  is called  $\ell$ -cosparse with respect to  $\Omega$ , if there is an index set  $\Lambda \subset [K]$  with  $|\Lambda| = \ell$ , such that  $\Omega_{\Lambda}y = 0$ . The support of a vector  $x \in \mathbb{R}^K$  is defined by  $\operatorname{supp}(x) = \{k \in [K] : x_k \neq 0\}$  and the cosupport accordingly as  $\operatorname{cosupp}(x) = \{k \in [K] : x_k = 0\}$ . Note that by definition we have  $\operatorname{supp}(x) \cup \operatorname{cosupp}(x) = [K]$ . For the

Both authors are with the Department of Mathematics, University of Innsbruck, Technikerstraße 13, 6020 Innsbruck, Austria. Email: michael.sandbichler@uibk.ac.at, karin.schnass@uibk.ac.at

runtime complexity R(n), we adopt standard Landau notation, i.e.  $R(n) = \mathcal{O}(f(n))$  means, there is a constant C > 0, such that for large n, the runtime R(n) satisfies  $R(n) \leq Cf(n)$ . Finally, the Frobenius norm of a matrix A is defined by  $||A||_F^2 := \operatorname{tr}(A^*A)$ .

# II. THE SEQUENTIAL ANALYSIS OPERATOR LEARNING ALGORITHM - SAOL

Since optimisation principles have already successfully led to sequential algorithms for dictionary learning, [15], [16], we will start our quest for a sequential algorithm by motivating a suitable optimisation principle for analysis operator learning. Suppose, we are given signals  $y_n \in \mathbb{R}^d$  that are perfectly cosparse in an operator  $\Omega$ , i.e.  $\Omega y_n$  has  $\ell$  zero entries or equivalently  $\Omega y_n - x_n = 0$  for some  $x_n$  which has  $K - \ell$  non zero entries. If we collect the signals  $y_n$ as columns in the matrix  $Y = (y_1 \dots y_N)$ , then by construction we have  $\Omega Y - X = 0$  for some  $X \in \mathcal{X}_\ell$  with  $\mathcal{X}_\ell := \{(x_1, x_2, \dots, x_N) \in \mathbb{R}^{K \times N} : |\operatorname{supp}(x_n)| = K - \ell\}$ . In the more realistic scenario, where the signals are not perfectly cosparse, we should still have  $\Omega Y - X \approx 0$ , which naturally leads to the following minimisation program to recover  $\Omega$ ,

$$\underset{\Gamma \in \mathcal{A}, X \in \mathcal{X}_{\ell}}{\arg\min} \|\Gamma Y - X\|_{F}^{2}.$$
 (1)

Apart from additional side constraints on  $\Gamma$ , such as incoherence, the optimisation program above has already been used successfully as starting point for the development of two analysis operator learning algorithms, Analysis K-SVD [14] and Analysis SimCO [2], [3]. AKSVD is an alternating minimisation algorithm, which alternates between finding the best  $X \in \mathcal{X}_{\ell}$  for the current  $\Gamma$  and updating  $\Gamma$  based on the current X. Since the update of  $\Gamma$  uses singular value decompositions, the computational complexity of the algorithm soon becomes intractable as d increases. ASimCO is a gradient descent algorithm with line search. It produces results similar to AKSVD and has the advantage that it does so with a fraction of the computational cost. Still, at closer inspection we see that the algorithm has two problematic aspects. First, the line search cannot be realised resource efficiently, since in each step several evaluations of the target function are necessary, which take up a lot of computation time. Moreover for each of these function evaluation we must either reuse the training data, thus incurring high storage costs, or use a new batch of data, thus needing a huge amount of training samples. The second problematic aspect of the ASimCO algorithm is that the estimated operators always stay within the manifold A. This might seem reasonable, as we can only get feasible solutions, but has the disadvantage that, even with the optimal step size, it can a take a large number of iterations to follow the steepest descent path. Still, if we consider the speed up of ASimCO with respect to AKSVD we see that gradient descent is a promising approach if we can avoid the line search and its associated problems.

To see that a gradient descent algorithm for our problem can also be sequential, let us rewrite our target function,  $g_N(\Gamma) = \min_{X \in \mathcal{X}_{\ell}} \|\Gamma Y - X\|_F^2$ . Abbreviating  $\Lambda_n = \operatorname{supp}(x_n)$  and  $\Lambda_n^c = \operatorname{cosupp}(x_n)$ , we have

$$g_{N}(\Gamma) = \sum_{n=1}^{N} \min_{x_{n}:|\Lambda_{n}|=K-\ell} \|\Gamma y_{n} - x_{n}\|_{2}^{2} =$$
  
= 
$$\sum_{n=1}^{N} \min_{x_{n}:|\Lambda_{n}|=K-\ell} (\|\Gamma_{\Lambda_{n}^{c}} y_{n}\|_{2}^{2} + \underbrace{\|\Gamma_{\Lambda_{n}} y_{n} - x_{n}\|_{2}^{2}}_{=0})$$
  
= 
$$\sum_{n=1}^{N} \min_{|J|=\ell} \|\Gamma_{J} y_{n}\|_{2}^{2} =: f_{N}(\Gamma).$$

Since the gradient of a sum of functions is the sum of the gradients of these functions, from  $f_N$  we see that the gradient of our objective function can be calculated sequentially. Before going into more details about how to avoid a line search and thus stay sequential, let us lose a few words about the uniqueness of the minima of our objective function.

If the signals are perfectly cosparse in  $\Omega$ , clearly there is a global minimum of  $f_N$  at  $\Omega$ . However, one can easily see that all permutations and sign flips of rows of  $\Omega$  are also minimisers of  $f_N$ . We call these the *trivial ambiguities*. The more interesting question is whether there are other global or local minima?

This question can be readily answered with an example. If all our training signals are (perfectly)  $\ell$ -cosparse in  $\Omega$  but lie in a subspace of  $\mathbb{R}^d$ , this will, at least without any further modifications, be problematic. In this case we can choose a vector v with  $||v||_2 = 1$  in the orthogonal complement of this subspace, and construct a continuum of operators  $\Gamma$ , which also satisfy  $f_N(\Gamma) = 0$ , by setting  $\gamma_k = a_k \omega_k + b_k v$  for some  $a_k^2 + b_k^2 = 1$ . This example indicates that isotropy in the data is important for our problem to be well posed. On the other hand, in case the data has such a low dimensional structure, which can be found via a singular value decomposition of  $Y^*Y$ , it is easy to transform the ill posed problem into a well posed one. Armed with the nonzero singular vectors, we just have to project our data onto the lower dimensional space spanned by these vectors and learn the analysis operator within this lower dimensional space. For simplicity we will from now on assume that any such preprocessing has already been done and that the data isotropically occupies the full ambient space  $\mathbb{R}^d$ or equivalently that  $Y^*Y$  is well conditioned.

# A. Minimising $f_N$

As mentioned above in order to get a sequential algorithm we want to use gradient descent but avoid the line search. Our strategy will be to use projected stochastic gradient descent with carefully chosen stepsize. Given the current estimate of the analysis operator  $\Gamma$ , one step of (standard) gradient descent takes the form

$$\bar{\Gamma} = \Gamma - \alpha \nabla f_N \left( \Gamma \right).$$

Let us calculate the gradient  $\nabla f_N(\Gamma)$  explicitly.<sup>1</sup> Denote by  $J_n$  the set for which  $\|\Gamma_{J_n} y_n\|_2^2 = \min_{|J|=\ell} \|\Gamma_J y_n\|_2^2$ , then the

<sup>&</sup>lt;sup>1</sup>The careful reader might observe that the function  $f_N$  is actually not differentiable everywhere. However this region is just a small set of measure zero. The computation given here is hence correct almost everywhere.

derivative of  $f_N$  with respect to a row  $\gamma_k$  of  $\Gamma$  is

$$\frac{\partial f_N}{\partial \gamma_k}(\Gamma) = \sum_{n=1}^N \sum_{j \in J_n} \frac{\partial}{\partial \gamma_k} \langle \gamma_j, y_n \rangle^2 =$$
$$= \sum_{n=1}^N \sum_{j \in J_n} 2\langle \gamma_j, y_n \rangle y_n^* \delta_{kj} = \sum_{n: \ k \in J_n} 2\langle \gamma_k, y_n \rangle y_n^* =: G_k.$$

Note that as expected the vectors  $G_k$  can be calculated sequentially, that is given a continuous stream of data  $y_n$ , we compute  $J_n$ , update all  $G_k$  for  $k \in J_n$ , and forget the existence of  $y_n$ . After processing all signals we set

$$\bar{\gamma}_k = (\gamma_k - \alpha_k G_k) \,\beta_k. \tag{2}$$

where  $\beta_k = \|\gamma_k - \alpha_k G_k\|_2^{-1}$  is a factor ensuring normalisation of  $\bar{\gamma}_k$ . This additional projection is necessary, since a standard gradient step will most likely take us out of the oblique manifold. If we compare to dictionary learning, e.g. [15], it is interesting to observe that we cannot simply choose  $\alpha_k$  by solving the linearised optimisation problem with side constraints using Lagrange multipliers since this would lead to a zero-update  $\bar{\gamma}_k = 0$ . Instead, we have to carefully choose the descent parameters  $\alpha_k$  in order to succeed. Since we want to find a minimum, the first idea is to select  $\alpha_k$  in a way, that we reduce our objective function in every step, meaning  $f_N(\bar{\Gamma}) \leq f_N(\Gamma)$ , as is ensured by the line search. To stay sequential and computationally light, we proceed along different lines. Observe that a decrease of the objective function is ensured if for all  $n \in [N]$  we have

$$\langle \bar{\gamma}_k, y_n \rangle^2 = \langle (\gamma_k - \alpha_k G_k) \beta_k, y_n \rangle^2 \le \langle \gamma_k, y_n \rangle^2.$$
 (3)

A simple calculation to be found in Appendix A shows that the condition on  $\alpha_k$  above is equivalent to

$$\alpha_{k} \leq \frac{2\langle G_{k}, y_{n} \rangle \langle \gamma_{k}, y_{n} \rangle - 2\langle \gamma_{k}, y_{n} \rangle^{2} \langle G_{k}, \gamma_{k} \rangle}{\langle G_{k}, y_{n} \rangle^{2} - \langle \gamma_{k}, y_{n} \rangle^{2} \langle G_{k}, G_{k} \rangle}, \quad (4)$$

making

$$\alpha_k := \min_n \frac{2\langle G_k, y_n \rangle \langle \gamma_k, y_n \rangle - 2\langle \gamma_k, y_n \rangle^2 \langle G_k, \gamma_k \rangle}{\langle G_k, y_n \rangle^2 - \langle \gamma_k, y_n \rangle^2 \langle G_k, G_k \rangle}$$
(5)

a viable choice.



Fig. 1. Decay of the target function after random initialisation for various choices of the stepsize, minimum vs. median with  $\alpha = 0.1$  (left) and median with varying prefactors  $\alpha$  (right). The 10000 noiseless, 92-cosparse training signals used per iteration were constructed according to the setup described in Section II-B with  $\Omega$  the 200  $\times$  100 Dirac-DCT operator.

Unfortunately, as can be seen in Figure 1(left), this choice typically yields way too small stepsizes because of possible outliers in the data. The small stepsizes in turn lead to a slow decrease of the objective function and slow convergence of the algorithm. Therefore we will lower our expectations and be satisfied, if we step towards lower values on average. This suggests as suitable choice

$$\alpha_{k} = \alpha \operatorname{median}_{n} \frac{2\langle G_{k}, y_{n} \rangle \langle \gamma_{k}, y_{n} \rangle - 2\langle \gamma_{k}, y_{n} \rangle^{2} \langle G_{k}, \gamma_{k} \rangle}{\langle G_{k}, y_{n} \rangle^{2} - \langle \gamma_{k}, y_{n} \rangle^{2} \langle G_{k}, G_{k} \rangle}.$$
(6)

The median has the advantage of being stable with respect to outliers and yielding descent for at least half of the  $y_n$ . The scaling factor  $\alpha \in (0,1)$  ensures that we have a descent for more than half of the  $y_n$  and therefore decrease the target function value. The disadvantage of choosing the stepsize according to (6) is that it cannot be done sequentially, since it presumes the knowledge of  $G_k$ , which is only available after processing all N signals of the current iteration. This would make it necessary to store all N signals  $y_n$  just to compute the descent parameter, which is certainly not viable in online learning. However, observe that  $G_k$  is itself an empirical estimator of  $\mathbb{E}_y \chi_{\{y: k \in J_y\}} 2\langle \gamma_k, y \rangle y^*$ , where  $J_y = \arg \min_{|J|=\ell} \|\Gamma_J y\|_2^2$  and hence  $G_k$  can be computed sequentially. We will use this fact and the stability of the median to construct the descent parameters sequentially. We first compute an approximation  $G_k$  to the gradient  $G_k$  using 95% or more generally  $N - L = N(1 - \varepsilon)$  of the total N signals. For the remaining signals in addition to updating  $G_k$ , we compute and store the  $L = \varepsilon N$  quantities

$$\alpha_k^{(n)} = \frac{2\langle G_k, y_n \rangle \langle \gamma_k, y_n \rangle - 2\langle \gamma_k, y_n \rangle^2 \langle G_k, \gamma_k \rangle}{\langle \tilde{G}_k, y_n \rangle^2 - \langle \gamma_k, y_n \rangle^2 \langle \tilde{G}_k, \tilde{G}_k \rangle}.$$
 (7)

Once all N signals have been processed, we compute the median of the  $\alpha_k^{(n)}$ , scale it with a fixed prefactor  $\alpha < 1$  and perform the gradient step. The choice of the prefactor is quite delicate since too small  $\alpha$  will decrease the learning rate of the algorithm while too large  $\alpha$  increase the risk of not decreasing the target value at all. Experiments suggest to choose  $\alpha$  between 0.1 and 0.3. From Figure 1 (right) we can see that for larger sizes of  $\alpha$  the algorithm decreases the target function sooner, but in the end does not reach small values. For very small  $\alpha$  the 500 iterations that were performed were not enough to decrease the target function sufficiently.

We summarise the first version of the algorithm, which takes as input parameters the current estimate of the analysis operator  $\Gamma \in \mathbb{R}^{K \times d}$ , the cosparsity parameter  $\ell$ , N training signals  $Y = (y_1, y_2, \dots, y_N)$ , the scaling factor  $\alpha$  and the parameter  $\varepsilon$  in Table I. Looking at the algorithm, we see that the first computationally expensive task is determining the sets  $J_n$ . This has to be done for each of our N sample vectors via determining the  $\ell$  smallest entries in the product  $\Gamma y_n$ . The matrix-vector product takes (2d-1)K operations, searching can be done in one run through the K resulting entries, yielding an overall runtime complexity of  $\mathcal{O}(dKN)$  for this step. The next expensive step is determining the step size  $\alpha$ . Each of the  $\varepsilon N$  fractions can be evaluated with  $\mathcal{O}(dN\varepsilon)$ operations, finding the median via quicksort can on average be done in  $\mathcal{O}(\varepsilon N \log(\varepsilon N))$ . Performing the gradient descent step is now cheap, as all of the used expressions have already been calculated. Overall, performing the aforementioned tasks





for k = 1, ..., K, for this step, we get a runtime complexity of  $\mathcal{O}(dKN + \varepsilon dN \log(\varepsilon N))) = \mathcal{O}(dKN)$  per iteration. Note that the storage requirements are only  $\mathcal{O}(dK)$ , amounting to storing the current iterate of the operator  $\Gamma$ . To see how the algorithm performs we will next conduct some experiments both on synthetic and image data.

#### B. Experiments on synthetic data

In the first set of experiments<sup>2</sup>, we use synthetic data generated from a given (target) analysis operator  $\Omega$ . A data vector yis generated by choosing a vector z from the unit sphere and a random subset  $\Lambda$  of  $\ell$  analysers. We then project z onto the orthogonal complement of the chosen analysers, contaminate it with Gaussian noise and normalise it, see Table II. The cosparse signals generated according to this model are very isotropic and thus do not exhibit the pathologies we described in the counterexample at the beginning of the section.

**Target operator:** As target operator for our experiments with synthetic data, we used a Dirac-DCT operator of size  $200 \times 100$  consisting of the identity matrix in the upper 100 rows and the DCT basis in the lower 100 rows. For illustration Figure 3 shows the Dirac-DCT operator of size  $40 \times 20$ .

<sup>2</sup>All experiments can be reproduced using the SAOL Matlab toolbox available at http://homepage.uibk.ac.at/~c7021041/code/SAOL.zip.

Signal model( $\Omega, \ell, \rho$ ) Input: •  $\Omega \in \mathbb{R}^{K \times d}$  - target analysis Operator, •  $\ell$  - cosparsity level of the signals w.r.t.  $\Omega$ , •  $\rho$  - noise level. Generation of the signals is done in the following way: • Draw  $z \sim \mathcal{N}(0, I_d), r \sim \mathcal{N}(0, \rho^2 I_d)$  and  $\Lambda \sim \mathcal{U}(\binom{[K]}{\ell})$ . • Set  $y = \frac{(\mathbb{1} - \Omega_{\Lambda}^{\dagger} \Omega_{\Lambda})z + r}{\|(\mathbb{1} - \Omega_{\Lambda}^{\dagger} \Omega_{\Lambda})z + r\|}.$  (8) The matrix  $(\mathbb{1} - \Omega_{\Lambda}^{\dagger} \Omega_{\Lambda})$  is a projector onto the space of all cosparse signals with cosupport  $\Lambda$ , so generating our signals in this way makes sure that they are (up to some noise) cosparse.

> TABLE II Signal model

**Training signals:** In each iteration of the algorithm, we use  $2^{14} = 16384$  signals drawn according to the signal model in Table II with cosparsity level  $\ell \in \{80, 88, 96\}$  and noiselevel  $\rho = 0$  resp.  $\rho = 0.2/\sqrt{d}$ .

**Initialisation & setup:** We use both a closeby and a random initialisation of the correct size. For the closeby initialisation, we mix the Dirac-DCT operator 1:1 with a random operator and normalise the rows, that is our initialisation operator is given by  $\Gamma_0 = D_n(\Omega + R)$ , where R is a  $K \times d$  matrix with rows drawn uniformly at random from the unit sphere  $\mathbb{S}^{d-1}$  and  $D_n$  is a diagonal matrix that ensures that the rows of  $\Gamma_0$  are normalised. For the random initialisation we simply set  $\Gamma_0 = R$ . The correct cosparsity level  $\ell$  is given to the algorithm and, as suggested by Figure 1, a prefactor  $\alpha = 0.1$  is chosen for the median, which is calculated from 5% of the signals,  $\varepsilon = 0.05$ . The results have been averaged over 20 runs with different initialisations.

**Recovery threshold:** We use the convention that an analyser  $\omega_k$  is recovered if  $\max_j |\langle \omega_k, \gamma_j \rangle| \ge 0.99$ .

The results of our first experiment are shown in Figure 2.



Fig. 2. Recovery rates of SAOL from signals with various cosparsity levels  $\ell$  in a noiseless (top) and a noisy setting (bottom), using a random (left) and a closeby (right) initialisation.

For the closeby initialisation we get (almost) perfect recovery of the target operator for all cosparsity levels, both in the noisy and the noiseless setting, which indicates that locally our algorithm performs as expected. For the random initialisation the algorithm tends to get saturate at around 70% recovery for all parameter settings. This is not surprising, as the nonconvex optimisation we perform depends heavily on the initialisation. In case of the closeby initialisation, we automatically set each of the rows of the starting operator near the desired row in the target operator. In contrast for the random initialisation it is very likely that two rows of the initialised operator lie close to the same row of the target operator. Our gradient descent algorithm then tends to find the nearest row of the target operator and thus we get multiple recovery of the same row. As we have prescribed a fixed number of rows, another row must be left out, which leads to the observed stagnation of the recovery rates and means that we are trapped in a local minimum of our target function. Figure 3 illustrates this effect for the Dirac-DCT operator in  $\mathbb{R}^{40\times 20}$ . Since the phenomenon of recovering duplicates is not only



Fig. 3. Operator learned with SAOL from a random initialisation (left) vs the original Dirac-DCT operator (right). The rows of the learned operator have been reordered and the signs have been matched with the original operator for easier comparison. For the learning 200 iterations with 8192 noiseless 12-cosparse signals, constructed according to the model in Table II, and a prefactor  $\alpha = 0.1$  were used.

as old as analysis operator learning but as old as dictionary learning, [1], there is also a known solution to the problem, which is the replacement of coherent analysers or atoms.

## C. Replacement

A straightforward way to avoid learning analysis operators with duplicate rows is to check after each iteration, whether two analysers of our current iterate  $\Gamma$  are very coherent. Under the assumption that the coherence of the target operator  $\mu(\Omega) = \max_{i \neq j \in [K]} |\langle \omega_i, \omega_j \rangle|$  is smaller than some threshold  $\mu(\Omega) \leq \mu_0$ , we know that two rows of  $\gamma_i, \gamma_j$  are likely to converge to the same target analyser, whenever we have  $|\langle \gamma_i, \gamma_j \rangle| > \mu_0$ . In this case, we discard one of two analysers, redraw it uniformly at random from the unit sphere  $\mathbb{S}^{d-1}$  and continue with the next iteration. Since, unlike dictionaries, analysis operators can be quite coherent and still perform very well, it is recommendable to be conservative and set the coherence threshold  $\mu_0$  rather high, in the extreme case as high as the recovery threshold 0.99.

Figure 4 shows the recovery results of our algorithm with



Fig. 4. Recovery rates of SAOL with replacement from signals with various cosparsity levels  $\ell$  in a noiseless (left) and a noisy setting (right), using a random initialisation.

the added replacement step for  $\mu_0 = 0.99$ , when using a

random initialisation and the same settings as described in Section II-B.

In the noiseless case, even when we employ the replacement strategy, we do not manage to get full recovery. If we introduce a small amount of noise, this problem vanishes and the algorithm returns the original operator. This is due to the dithering effect of the noise, which breaks symmetries in the synthetic data, and which we can assume to be present in all real life data, in particular the image data which we will use in our second set of experiments.

# D. Experiments on image data

To get an indication how our algorithm performs on real data, we will use it to learn a quadratic analysis operator on all  $8 \times 8$  patches of the  $256 \times 256$  Fabio image, cf. Figure 10. Since we do not have a reference dictionary for comparison this time, we look at the target function after each iteration. We initialise the analysis operator  $\Gamma \in \mathbb{R}^{64 \times 64}$  randomly as for the synthetic data and set the cosparsity level  $\ell = 57$ , the parameter  $\varepsilon = 0.05$  and the replacement threshold  $\mu_0 = 0.99$ . For each iteration we choose 15000 out the available 62001 patches of Fabio uniformly at random as training signals. Figure 5 shows the decay of the target function for several scaling factors  $\alpha$  as well as the learned operator for  $\alpha = 0.1$ .



Fig. 5. Decay of the target function using SAOL with various prefactors  $\alpha$  (left) for the Fabio image and the recovered operator for  $\alpha = 0.1$  (right).

Figure 5 (left) shows again how the size of the prefactor  $\alpha$ affects the stability of the algorithm. We can also see that since the stepsize has to be chosen rather small, training of the analysis operator takes a significant number of iterations, which is rather disappointing. Increasing the stepsize to speed up convergence has the opposite effect and results in failure to minimise the target function. Still, if we look at the learned operator for  $\alpha = 0.1$ , we can see the merit of our method. The operator seems to consist of pooled edge detectors, which are known to cosparsify grayscale images. Note also that the  $d \times d$  analysis operator is naturally very different from any  $d \times d$  dictionary we could have learned with corresponding sparsity level  $S = d - \ell$ , see e.g [16]. This is due to the fact that image patches are not isotropic, but have their energy concentrated in the low frequency ranges. So while both the  $d \times d$  dictionary and analysis operator will not have (stable) full rank, the dictionary atoms will tend to be in the low frequency ranges, and the analysers will - as can be seen - tend to be in the high frequency ranges.

We also want to mention that for image data the replacement strategy for  $\mu_0 = 0.99$  is hardly ever activated. Lowering the threshold results in continuous replacement and refinding of the same analysers. This phenomenon is again explained by the lack of isotropy and the shift invariant structure of the patch data, for which translated and thus coherent edge detectors, as seen in Figure 5, naturally provide good cosparsity.

Encouraged by the learned operator we will explore in the next section how to stabilise the algorithm and accelerate its convergence.

#### III. THE IMPLICIT SAOL ALGORITHM - ISAOL

Due to the stepsize problems on real data, we need to rethink our approach and try to enforce stability of the algorithm. In standard gradient descent, for each row of  $\Gamma$ , we have the iteration

$$\bar{\gamma}_k = \gamma_k - \alpha \nabla f_N(\Gamma)_k. \tag{9}$$

Rewriting yields

$$\frac{\bar{\gamma}_k - \gamma_k}{\alpha} = -\nabla f_N(\Gamma)_k, \tag{10}$$

which can be interpreted as an explicit Euler step for the system of ordinary differential equations

$$\dot{\gamma}_k = -\nabla f_N(\Gamma)_k, \ k \in [K]. \tag{11}$$

As it is the simplest integration scheme for ordinary differential equations, the explicit Euler scheme is known to have a very limited region of convergence with respect to the stepsize. In our case, this means that we have to choose extremely small values for the descent parameter  $\alpha$  in order to achieve convergence.

The tried and tested strategy to get rid of stability issues when numerically solving differential equations is to use an implicit scheme for the integration [8], [9]. We will use this as an inspiration to obtain a more stable learning algorithm.

We shortly recall the notion of an implicit integration scheme. Suppose we want to solve the differential equation  $\dot{x} = f(x)$ . If we discretise x(t) and approximate the derivative by  $\dot{x}(t_n) \approx \frac{x(t_n)-x(t_{n-1})}{t_n-t_{n-1}}$ , we have to choose whether we use the approximation  $\dot{x}(t_n) = f(x(t_n))$  or  $\dot{x}(t_n) = f(x(t_{n-1}))$ . Choosing  $f(x(t_{n-1}))$  yields the explicit Euler scheme, which in our setting corresponds to the SAOL algorithm. If we choose  $f(x(t_n))$  we obtain the implicit Euler scheme and need to solve

$$\frac{x(t_n) - x(t_{n-1})}{t_n - t_{n-1}} = f(x(t_n)).$$
(12)

If f(x) = Ax is linear, this leads to the recursion

$$x(t_n) = (\mathbb{1} - (t_n - t_{n-1})A)^{-1}x(t_{n-1}),$$
(13)

and in each step we need to solve a system of linear equations. This makes implicit integration schemes inherently more expensive than explicit schemes. However, in return we get additional stability with respect to the possible stepsizes. If fis a nonlinear function, the inversion is more difficult and can often only be approximated for example via a Newton method.  $\textbf{ISAOL}(\Gamma,\ell,Y)$  - (one iteration)



Mapping everything to our setting, we observe that the gradient  $\nabla f_N(\Gamma)$  is nonlinear because the sets  $J_n$  depend on  $\Gamma$ . Still, due to the special structure of the gradient  $\nabla f_N(\Gamma)$ , it has a simple linearisation,  $\nabla f_N(\Gamma)_k = 2\gamma_k \sum_{n: k \in J_n} y_n y_n^*$ . We can now use the current iterate of  $\Gamma$  to compute the matrix  $A_k(\Gamma) := \sum_{n: k \in J_n} y_n y_n^*$  and to linearise the equation. For our operator learning problem, we get the following linearised variant of the implicit Euler scheme

$$\frac{\gamma_k - \gamma_k}{\alpha} = -\bar{\gamma}_k A_k(\Gamma), \tag{14}$$

leading to the recursion

$$\bar{\gamma}_k = \gamma_k (\mathbb{1} + \alpha A_k(\Gamma))^{-1}.$$
(15)

Due to the unconditional stability of the implicit Euler scheme [9], we can take  $\alpha$  constant 1. In order to stay within the oblique manifold, we again perform a projection to the unit sphere after each step. The final algorithm is summarised in Table III.

Let us again take a short look at the computational complexity of the implicit algorithm and the price we have to pay for increased stability. As in the previous section, we need to compute all products of the vectors  $y_n$  with the current iterate  $\Gamma$ , cost  $\mathcal{O}(NKd)$ . Furthermore, in each step we need to solve K linear systems of size  $d \times d$  amounting to an additional cost of  $\mathcal{O}(Kd^2)$ . So altogether for one step, we arrive at  $\mathcal{O}(NKd+Kd^2) = \mathcal{O}(NKd)$ . However, in constrast to SAOL, if we want to preserve sequentiality of the algorithm we have to store the K matrices  $A_k$  in each step, amounting to an additional spatial complexity of  $\mathcal{O}(Kd^2)$ . In a non-sequential setting and in case N < Kd this can be reduced to the storage cost of the data matrix  $\mathcal{O}(Nd)$ .

#### A. Experiments on synthetic data

As in the previous section, we first try our new algorithm on synthetic data. For this, we again learn an operator from data generated with the Dirac-DCT operator. The setup is the same as in Section II-B and the results are shown in Figure 6. Note that the recovery with ISAOL is slightly slower than in the case of SAOL, Figure 2. This happens because for synthetic data the SAOL algorithm is able to choose fairly large stepsizes. The ISAOL algorithm is more pessimistic and takes slightly longer to converge. In the next section we will see that this trend is drastically reversed for image data, where the stepsize chosen by the SAOL algorithm tends be much smaller.

Finally, since the implicit algorithm per se, like SAOL, does not penalise the recovery of two identical rows, cf. Figure 3,

7



Fig. 6. Recovery rates of ISAOL from signals with various cosparsity levels  $\ell$  in a noisy setting, using a random (left) and a closeby (right) initialisation.

we need to again use the replacement strategy introduced in Section II-C. The simulation results, using replacement with



Fig. 7. Recovery rates of ISAOL with replacement from signals with various cosparsity levels  $\ell$  in a noiseless (left) and a noisy setting (right), using a random initialisation.

 $\mu_0 = 0.99$  and the usual setup are shown in Figure 7. We see that contrary to the SAOL algorithm the ISAOL algorithm achieves full recovery not only for noisy but also for noiseless data, see Figure 4. This comes, however, at the cost of a significant number of iterations.

#### B. Experiments on image data

Finally, we want to see how the stabilised algorithm performs on real data. We use the same image (Fabio) and setup as in Section II-D to learn a square analysis operator for  $8 \times 8$ patches, cf. Figure 5.

As can be seen in Figure 8, the training is much faster now,



Fig. 8. Decay of the target function using (I)SAOL for the Fabio image (left) and the operator recovered by ISAOL (right).

because the stepsize does not have to be chosen as fine as in the previous section. The decrease in the objective function is very fast compared to SAOL, and we see that already after a few iterations the algorithm stabilises and we obtain similar checkerboard-like structures as in Figure 5. As for SAOL we observe that the replacement strategy for  $\mu_0 = 0.99$  is hardly ever activated and that lowering the threshold results in finding and replacing the same translated edge detectors.

We are now ready to compare the performance of the developed algorithms to their closest counterparts Analysis SimCO (ASimCO) and Incoherent Analysis SimCO (IASimCO), where IASimCo is ASimCo with a replacement strategy, [3].

#### IV. (I)SAOL VS. (I)ASIMCO

We again conduct experiments on both synthetic and image data. For the synthetic data we use the same setup as described in Section II-B, that is we try to recover the Dirac-DCT operator in  $\mathbb{R}^{200\times100}$  from  $2^{14} = 16384$  signals drawn in each iteration according to the signal model in Table II with cosparsity level  $\ell = 96$  and noiselevel  $\rho = 0.2/\sqrt{d}$ . For SAOL we choose  $\alpha = 0.1$  and  $\varepsilon = 0.05$ . As replacement threshold for both algorithms we use  $\mu_0 = 0.8$  corresponding to the threshold of the out-of-the-box version of IASimCO, available on the authors' homepage, which was used here.

Figure 9 (left) shows the recovery rates of the 4 algorithms. We can see that ISAOL starts out fastest, followed by SAOL and IASimCo but is then overtaken first by SAOL and then IASimCO. The fact that after 1000 iterations IASimCO finishes as close second after SAOL, while ASimCO saturates at around 70% recovery suggests that for perfectly isotropic data the replacement strategy employed in IASimCO is more efficient than our random replacement strategy. When comparing the average calculation times per iteration on a 3.1 GHz Intel Core i7 Processor, we find that SAOL is about 2 times faster than (I)ASimCO, while ISAOL is about 1.5 times slower. In the experiment on image data, we learn an overcomplete



Fig. 9. Recovery rates of (I)ASimCo and (I)SAOL from 96-cosparse signals in a noisy setting (left). Decay of the target function using (I)ASimCo and (I)SAOL to learn a  $128 \times 64$  operator for the House image (right).

operator with 128 rows from the  $8 \times 8$  patches of the  $256 \times 256$ (unnormalised) House image contamined with Gaussian noise with  $\sigma = 12.8$ , corresponding to PSNR  $\approx 25$ . Motivated by the choice of parameters in [3] we choose as cosparsity level  $\ell = 50$ , initialise randomly and in each iteration use 20000 randomly selected patches out of the available 62001. As usual for SAOL we use  $\alpha = 0.1$  and  $\varepsilon = 0.05$ . Since for image data our replacement strategy is hardly ever activated, we directly omit it to save computation time. (I)ASimCo are again used in their out of the box versions. Figure 9 shows that the ISAOL algorithm indeed minimises the target function in a fraction of the iterations necessary for the SAOL algorithm, which in turn is much faster than (I)ASimCO. Already after 100 iterations the ISAOL algorithm has essentially finished minimising the objective function, whereas SAOL needs 1000 iterations to get to approximately the same value of the objective function. Both ASimCO and IASimCo lag behind



Fig. 10. Images used for learning and denoising. Top: Fabio, House, Peppers; Bottom: Cameraman, Barbara, Mandrill.

and, as indicated by the shape of the curve, would need much more than 1000 iterations to approach a comparable target function value. Finally note that the equal performance of ASimCO and IASimCO again indicates that for image data replacement strategies hardly make a difference.

Encouraged by this good performance we will in the next section apply both our algorithms to image denoising.

#### V. IMAGE DENOISING

In this section we will compare the performance of analysis operators learned by (I)SAOL in combination with Tikhonov regularisation for image denoising to the performance of operators learned by (I)ASimCO. For easy comparison we use the same setup as in [3], where (I)ASimCo is compared to several other major algorithms for analysis operator learning, [12], [19], [10], [14], [4], and found to give the best performance. **Learning setup:** We follow the setup for the House image in the last section. Our training data consists of all  $8 \times 8$  patches of one of the  $256 \times 256$  images from Figure 10 corrupted with

Gaussian white noise of level  $\sigma = 12.8$  resp.  $\sigma = 45$  leading to a PSNR of approximately 15dB resp. 25dB. The analysis operators of size  $128 \times 64$  are initialised by drawing each row uniformly at random from the unit sphere, and then updated using in each step 20000 randomly selected patches of the available 62001 and a cosparsity level  $\ell \in \{40, 50, 60, 70, 80\}$ . The same initialisation is used for all four algorithms. For (I)ASimCo and SAOL we use 2000 and for ISAOL 500 iterations. For SAOL we choose as usual  $\alpha = 0.1$  and  $\varepsilon = 0.05$ and again we omit the replacement step for (I)SAOL.

**Denoising setup:** For the denoising step we use a standard approach via Tikhonov regularisation based on the learned analysis operator  $\Omega$ , [5], [6]. For each noisy patch y we solve,

$$\hat{y} = \arg\min \,\lambda \|\Omega z\|_1 + \|z - y\|_2$$
 (16)

for a regularisation parameter  $\lambda \in \{0.002, 0.01, 0.05, 0.1, 0.3, 0.5\}$ . We then reassemble the denoised patches  $\hat{y}$  to the denoised image, by averaging each pixel in the full image over the denoised patches in which it is contained. To measure the quality of the reconstruction for each cosparsity level  $\ell$  and regularisation parameter  $\lambda$  we average the PSNR of the denoised image over 5 different noise realisations and initialisations. Table IV

shows the PSNR for optimal choice of  $\ell$  and  $\lambda$  for each of the algorithms. We can see that all four algorithms give a comparable denoising performance. In the lower noise regime (I)ASimCO has a slight advantage for the untextured images, Peppers and Cameraman, while for the textured images, House, Barbara and Mandrill, (I)SAOL provides better results. In the high noise regime (I)SAOL always performs slightly better. The optimal parameters are quite stable across images and algorithms and seem to only depend on the noise levels. In the lower noise regime the optimal parameter  $\lambda$  for (I)ASimCo is 0.1 or 0.3, while for (I)SAOL it is 0.05 or 0.1 In the high noise regime the optimal parameter  $\lambda$  for all four algorithms is 0.01. It is interesting to observe that (I)ASimCO always has the best performance for the highest co-sparsity level in the training  $\ell = 80$ . For (I)SAOL this is only true in the higher noise regime, where the zeros tend to be masked by the noise. In the lower noise regime the performance is stable for  $\ell$  between 60 and 80. For the interested reader we provide the denoising results on Barbara and Peppers for all parameter choices in the appendix, Figure 11.

After confirming that our algorithms indeed learn useful operators also on real data we now turn to a discussion of our results.

# VI. DISCUSSION

We have developed two algorithms for analysis operator learning based on projected stochastic gradient descent, SAOL and ISAOL. The algorithms perform slightly better than the state of the art algorithms (I)ASimCO, [3], which are similarly gradient descent based and have slightly higher but comparable computational complexity per iteration, in terms of recovery rates resp. reduction of the objective function. Another advantage of SAOL is that it is sequential with a memory requirement corresponding to the size of the operator,  $\mathcal{O}(dK)$ ). In contrast ASimCO either is non sequential with a memory requirement of the order of the data matrix,  $\mathcal{O}(dN)$ , or in a sequential setting needs O(LN) training sample corresponding to the L evaluations of the objective function necessary for the line search. ISAOL, which is more stable than SAOL, is sequential when accepting a memory requirement  $\mathcal{O}(d^2K)$ and in a non sequential setting has again memory requirement  $\mathcal{O}(dN).$ 

Considering image denoising via Tikhonov regularisation as application of analysis operator learning, we see that the (I)SAOL operators give slightly better results than the

Algorithm	$\sigma(\text{PSNR})$	Pep	Cam	Hou	Bar	Man
ASimCO	12.8(25)	31.29	30.35	30.99	30.49	28.31
IASimCO		31.32	30.31	30.59	30.33	28.05
SAOL		31.26	30.27	32.67	30.88	28.76
ISAOL		31.08	30.24	32.73	30.85	28.62
ASimCO	45(15)	24.53	23.31	26.80	25.23	23.07
IASimCO		24.32	23.05	26.34	25.13	22.91
SAOL		24.50	23.24	26.53	25.12	23.30
ISAOL		24.64	23.39	27.05	25.32	23.33

TABLE IV Performance of (I)ASimCO and (I)SAOL for denoising for different pictures and noise levels.

(I)ASimCo operators for textured images or high noise levels (up to 1.5dB) and slightly worse results for untextured images and low noise levels (at worst 0.1dB).

A Matlab toolbox to reproduce all the experiments reported in this paper can be found at http://homepage.uibk.ac.at/  $\sim$ c7021041/code/SAOL.zip.

While the good performance of the developed algorithms certainly justified the effort, one our main motivations for considering a projected gradient descent approach to analysis operator learning was to derive convergence results similar to those for dictionary learning, [16]. However, even a local convergence analysis, turns out to be quite different and much more complicated than for dictionary learning. The main reason for this is that sparsity is more robust to perturbations than co-sparsity. So for an S-sparse signal  $y = \Phi_I x_I$  and a perturbed dictionary  $\Psi$  with  $\|\psi_k - \phi_k\|_2 < \varepsilon$  for balanced  $x_I$ the best S-term approximation in  $\Psi$  will still use the same support I. In contrast if y is  $\ell$ -cosparse with respect to an analysis operator  $\Omega$ ,  $\Omega_{\Lambda} y = 0$ , then for a perturbed operator  $\Gamma$  with  $\|\gamma_k - \omega_k\|_2 < \varepsilon$  the smallest  $\ell$  entries of  $\Gamma y$  will not all be located in  $\Lambda$ . To get a local convergence result one has to deal with the fact that only part of the cosupport is preserved. We expect that for most signals containing k in the cosupport with respect to  $\Omega k$  will also be in the cosupport with respect to  $\Gamma$ . Unfortunately the mathematical tools necessary to quantify these statements are much more involved that the comparatively simple results necessary for the convergence of dictionary learning and so the local convergence analysis remains on our agenda for future research.

Another research direction, we are currently pursuing, is inspired by the shape of the analysis operators learned on noiseless images. The translation invariance of the edge detector like analysers suggests to directly assume translation invariance of the analysis operator. Such an operator has two advantages, first, learning it will require less training samples and second, since it can be reduced to several translated mother functions it will be cost efficient to store and apply.

#### **ACKNOWLEDGEMENTS**

This work was supported by the Austrian Science Fund (FWF) under Grant no. Y760. In addition the computational results presented have been achieved (in part) using the HPC infrastructure LEO of the University of Innsbruck. Part of this work has been carried out while M.S. was supported by the trimester program 'Mathematics of Signal Processing' at the Hausdorff Research Institute for Mathematics.

# APPENDIX A CONDITION ON $\alpha_k$

So satisfy the condition in (3) we need  $\langle \gamma_k - \alpha_k G_k, y_n \rangle^2 \leq \beta_k^{-2} \langle \gamma_k, y_n \rangle^2$  for  $\beta_k = \|\gamma_k - \alpha_k G_k\|_2^{-1}$ . Expanding the inner products we see that this is equivalent to,

$$\langle \gamma_k, y_n \rangle^2 - 2\alpha_k \langle \gamma_k, y_n \rangle \langle G_k, y_n \rangle + \alpha_k^2 \langle G_k, y_n \rangle^2 \leq \langle \gamma_k, y_n \rangle^2 (1 - 2\alpha_k \langle G_k, \gamma_k \rangle + \alpha_k^2 \langle G_k, G_k \rangle).$$



Fig. 11. PSNR for denoising of the Barbara and Peppers images for various cosparsity levels and regularisation parameters. From top to bottom: Barbara (PSNR 25), Peppers (PSNR 25), Barbara (PSNR 15), Peppers (PSNR 15).

Since we want  $\alpha_k \neq 0$ , this is equivalent to

$$-2\langle \gamma_k, y_n \rangle \langle G_k, y_n \rangle + \alpha_k \langle G_k, y_n \rangle^2 \leq -2\langle \gamma_k, y_n \rangle^2 \langle G_k, \gamma_k \rangle + \alpha_k \langle \gamma_k, y_n \rangle^2 \langle G_k, G_k \rangle,$$

so after doing the housekeeping we arrive at the condition,

$$\alpha_k \le 2 \frac{\langle \gamma_k, y_n \rangle \langle G_k, y_n \rangle - \langle \gamma_k, y_n \rangle^2 \langle G_k, \gamma_k \rangle}{\langle G_k, y_n \rangle^2 - \langle \gamma_k, y_n \rangle^2 \langle G_k, G_k \rangle}.$$

## Appendix B

# Some denoising results

We here provide the full denoising results for all used parameter choices  $\lambda \in \{0.002, 0.01, 0.05, 0.1, 0.3, 0.5\}$ , cosparsity levels  $\ell \in \{40, 50, 60, 70, 80\}$  and both noise levels for the images Barbara and Peppers, Figure 11.

#### REFERENCES

- M. Aharon, M. Elad, and A.M. Bruckstein. K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on Signal Processing.*, 54(11):4311–4322, November 2006.
- [2] J. Dong, W. Wang, and W. Dai. Analysis SimCO: A new algorithm for analysis dictionary learning. In ICASSP14, 2014.
- [3] J. Dong, W. Wang, W. Dai, M.D. Plumbley, Z. Han, and J. Chambers. Analysis SimCO algorithms for sparse analysis model based dictionary learning. *IEEE Transactions on Signal Processing*, 64(2):417–431, 2016.
- [4] E. M. Eksioglu and O. Bayir. K-SVD meets transform learning: Transform K-SVD. *IEEE Signal Processing Letters*, 21(3):347–351, 2014.
- [5] M. Elad and M. Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image* processing, 15(12):3736–3745, 2006.
- [6] M. Elad, R. Milanfar, and R. Rubinstein. Analysis versus synthesis in signal priors. *Inverse Problems*, 23(3):947–968, 2007.
- [7] R. Giryes, S. Nam, M. Elad, R. Gribonval, and M.E. Davies. Greedylike algorithms for the cosparse analysis model. *Linear Algebra and its Applications*, 441:22–60, 2014.
- [8] E. Hairer, S. P. Nørsett, and G. Wanner. Solving ordinary differential equations. I, volume 8 of Springer Series in Computational Mathematics. Springer-Verlag, Berlin, second edition, 1993. Nonstiff problems.
- [9] E. Hairer and G. Wanner. Solving ordinary differential equations. II, volume 14 of Springer Series in Computational Mathematics. Springer-Verlag, Berlin, 2010. Stiff and differential-algebraic problems, Second revised edition, paperback.
- [10] S. Hawe, M. Kleinsteuber, and K. Diepold. Analysis operator learning and its application to image reconstruction. *IEEE Transactions on Image Processing*, 22(6):2138–2150, 2013.

- [11] S. Nam, M.E. Davies, M. Elad, and R. Gribonval. The cosparse analysis model and algorithms. *Applied Computational Harmonic Analysis*, 34:30–56, 2013.
- [12] S. Ravishankar and Y. Bresler. Learning sparsifying transforms. *IEEE Transactions on Signal Processing*, 61(5):1072–1086, 2013.
- [13] R. Rubinstein, A. Bruckstein, and M. Elad. Dictionaries for sparse representation modeling. *Proceedings of the IEEE*, 98(6):1045–1057, 2010.
- [14] R. Rubinstein, T. Peleg, and M. Elad. Analysis K-SVD: A dictionarylearning algorithm for the analysis sparse model. *IEEE Transactions on Signal Processing*, 61(3):661–677, 2013.
- [15] K. Schnass. Local identification of overcomplete dictionaries. *Journal of Machine Learning Research (arXiv:1401.6354)*, 16(Jun):1211–1242, 2015.
- [16] K. Schnass. Convergence radius and sample complexity of ITKM algorithms for dictionary learning. accepted to Applied and Computational Harmonic Analysis (arXiv:1503.07027), 2016.
- [17] M. Seibert, J. Wörmann, R. Gribonval, and M. Kleinsteuber. Learning co-sparse analysis operators with separable structures. *IEEE Transactions on Signal Processing*, 64(1):120–130, 2016.
- [18] M. Yaghoobi, S. Nam, R. Gribonval, and M.E. Davies. Analysis operator learning for overcomplete cosparse representations. In *EUSIPCO11*, pages 1470–1474, 2011.
- [19] M. Yaghoobi, S. Nam, R. Gribonval, and M.E. Davies. Constrained overcomplete analysis operator learning for cosparse signal modelling. *IEEE Transactions on Signal Processing*, 61(9):2341–2355, 2013.