



Published in final edited form as:

IEEE Signal Process Lett. 2019 August ; 26(8): 1137–1141. doi:10.1109/lsp.2019.2921446.

Convolutional Analysis Operator Learning: Dependence on Training Data

Il Yong Chun[†] [Member, IEEE], David Hong[†] [Student Member, IEEE]

Department of Electrical Engineering and Computer Science, The University of Michigan, Ann Arbor, MI 48019 USA

Ben Adcock,

Department of Mathematics, Simon Fraser University, Burnaby, BC V5A 1S6 Canada

Jeffrey A. Fessler [Fellow, IEEE]

Department of Electrical Engineering and Computer Science, The University of Michigan, Ann Arbor, MI 48019 USA

Abstract

Convolutional analysis operator learning (CAOL) enables the unsupervised training of (hierarchical) convolutional sparsifying operators or autoencoders from large datasets. One can use many training images for CAOL, but a precise understanding of the impact of doing so has remained an open question. This paper presents a series of results that lend insight into the impact of dataset size on the filter update in CAOL. The first result is a general deterministic bound on errors in the estimated filters, and is followed by a bound on the *expected* errors as the number of training samples increases. The second result provides a *high probability* analogue. The bounds depend on properties of the training data, and we investigate their empirical values with real data. Taken together, these results provide evidence for the potential benefit of using more training data in CAOL.

I. Introduction

LEARNING convolutional operators from large datasets is a growing trend in signal/image processing, computer vision, machine learning, and artificial intelligence. The *convolutional* approach resolves the large memory demands of patch-based operator learning and enables unsupervised operator learning from “big data,” i.e., many high-dimensional signals. See [1], [2] and references therein. Examples include *convolutional dictionary learning* [2], [3] and *convolutional analysis operator learning* (CAOL) [1], [4]. CAOL trains an autoencoding CNN in an unsupervised manner, and is useful for training multi-layer CNNs from many training images [1]. In particular, the block proximal gradient method using a majorizer [1], [2] leads to rapidly converging and memory-efficient CAOL [1]. However, a theoretical

Personal use is permitted, but republication/redistribution requires IEEE permission. See <http://www.ieee.org/publicationsstandards/publications/rights/index.html> for more information.

iyunchun@umich.edu.

[†]The first two authors contributed equally to this work.

understanding of the impact of using many training images in CAOL has remained an open question.

This paper presents new insights on this topic. Our first main result provides a *deterministic* bound on filter estimation error, and is followed by a bound on the *expected* error when “model mismatch” has zero mean. (See Theorem 1 and Corollary 2, respectively.) The expected error bound depends on the training data, and we provide empirical evidence of its decrease with an increase in training samples. Our second main result provides a *high probability* bound that explicitly decreases with increasingly many i.i.d. training samples. The bound improves when model mismatch and samples are uncorrelated. (See Theorem 3.) Additional empirical findings provide evidence that the correlation can indeed be small in practice. Put together, our findings provide new insight into how using many samples can improve CAOL, underscoring the benefits of the low memory usage of CAOL.

II. Backgrounds and Preliminaries

A. CAOL with orthogonality constraints

CAOL seeks a set of filters that “best” sparsify a set of training images $\{x_l \in \mathbb{C}^N : l = 1, \dots, L\}$ by solving the optimization problem [1, §II-A] (see Appendix for notation):

$$\underset{D = [d_1, \dots, d_K]}{\operatorname{argmin}} \min_{\{z_{l,k}\}} F(D, \{z_{l,k}\}), \text{ subj. to } DD^H = \frac{1}{R} \cdot I, \quad (\text{P0})$$

$$F(D, \{z_{l,k}\}) := \sum_{l=1}^L \sum_{k=1}^K \|d_k \otimes x_l - z_{l,k}\|_2^2 + \alpha \|z_{l,k}\|_0$$

where \otimes denotes convolution, $\{d_k \in \mathbb{C}^R : k = 1, \dots, K\}$ is a set of K R convolutional kernels, $\{z_{l,k} \in \mathbb{C}^N : l = 1, \dots, L, k = 1, \dots, K\}$ is a set of sparse codes, $\alpha > 0$ is a regularization parameter controlling the sparsity of features $\{z_{l,k}\}$, and $\|\cdot\|_0$ denotes the ℓ_0 -quasi-norm. We group the K filters into a matrix:

$$D := [d_1 \cdots d_K] \in \mathbb{C}^{R \times K}. \quad (1)$$

The orthogonality condition $DD^H = \frac{1}{R}I$ in (P0) enforces 1) a tight-frame condition on the filters, i.e., $\sum_{k=1}^K \|d_k \otimes x\|_2^2 = \|x\|_2^2, \forall x$ [1, Prop. 2.1]; and 2) filter diversity when $R = K$, since $DD^H = \frac{1}{R}I$ implies $D^H D = \frac{1}{K}I$ and each pair of filters is incoherent, i.e., $|\langle d_k, d_{k'} \rangle|^2 = 0, \forall k \neq k'$. One often solves (P0) iteratively, by alternating between optimizing D (filter update) and optimizing $\{z_{l,k} : \forall l, k\}$ (sparse code update) [1], i.e., at the i iteration, the current iterates are updated as $\{z_{l,k}^{(i+1)}\} = \operatorname{argmin}_{\{z_{l,k}\}} F(D^{(i)}, \{z_{l,k}\})$ and $D^{(i+1)} = \operatorname{argmin}_{DD^H = \frac{1}{R} \cdot I} F(D, \{z_{l,k}^{(i+1)}\})$.

B. Filter update in a matrix form

The key to our analysis lies in rewriting the filter update for (P0) in matrix form, to which we apply matrix perturbation and concentration inequalities. Observe first that

$$d_k \otimes x_l = \underbrace{\left[\Pi^0 x_l, \dots, \Pi^{R-1} x_l \right]}_{=: \Psi_l \in \mathbb{C}^{N \times R}} d_k = \Psi_l d_k, \quad l = 1, \dots, L, \quad (2)$$

where $\Pi := \begin{bmatrix} 0 & I_{N-1} \\ 1 & 0 \end{bmatrix} \in \mathbb{C}^{N \times N}$ is the circular shift operator and $(\cdot)^n$ denotes the matrix product of its n copies. We consider a circular boundary condition to simplify the presentation of $\{\Psi_l\}$ in (2), but our entire analysis holds for a general boundary condition with only minor modifications of $\{\Psi_l\}$ as done in [1, §IV-A]. Using (2), the filter update of (P0) is rewritten as

$$D^* = \underset{D}{\operatorname{argmin}} \sum_{l=1}^L \|\Psi_l D - Z_l\|_F^2, \text{ subj. to } DD^H = \frac{1}{R} \cdot I, \quad (P1)$$

where $Z_l := [z_{l,1}, \dots, z_{l,K}] \in \mathbb{C}^{N \times K}$ contains all the current sparse code estimates for the l th sample, and we drop iteration superscript indices $(\cdot)^{(i)}$ throughout. The next section uses this form to characterize the filter update solution D^* .

III. Main Results:

Dependence of CAOL on Training Data

The main results in this section illustrate how training with many samples can reduce errors in the filter D^* from (P1) and characterize the reduction in terms of properties of the training data. Throughout we model the current sparse codes estimates as

$$Z_l = \underbrace{\Psi_l D_{\text{true}}}_{=: Z_{\text{true},l}} + E_l, \quad l = 1, \dots, L \quad (3)$$

where D_{true} is formed from optimal (orthogonal) filters analogously to (1), and $E_l \in \mathbb{C}^{N \times K}$ captures model mismatch in the current sparse codes, e.g., due to the current iterate being far from convergence or being trapped in local minima.

The following theorem provides a *deterministic* characterization.

Theorem 1.—Suppose that both matrices

$$\sum_{l=1}^L \Psi_l^H Z_l \in \mathbb{C}^{R \times K} \text{ and } \sum_{l=1}^L \Psi_l^H Z_{\text{true},l} \in \mathbb{C}^{R \times K} \quad (4)$$

are full row rank, where $\{\Psi_l, Z_l, Z_{\text{true},l} : l = 1, \dots, L\}$ are defined in (2)–(3). Then, the solution D^* to (P1) has error with respect to D_{true} bounded as

$$\|D^* - D_{\text{true}}\|_F^2 \leq 5 \frac{\|\sum_{l=1}^L \Psi_l^H E_l\|_F^2}{\lambda_{\min}^2(\sum_{l=1}^L \Psi_l^H \Psi_l)}, \quad (5)$$

where $\lambda_{\min}(\cdot)$ denotes the smallest eigenvalue of its argument.

The full row rank condition on (4) ensures that the estimated filters D^* and the true filters D_{true} are unique, and it further guarantees that the denominator of (5) is strictly positive. When the model mismatches E_1, \dots, E_L are independent and mean zero, we obtain the following expected error bound:

Corollary 2.—Under the construction of Theorem 1, suppose that E_l is a zero-mean random matrix for $l = 1, \dots, L$, and is independent over l . Then,

$$\mathbb{E}\|D^* - D_{\text{true}}\|_F^2 \leq 5\bar{\sigma}^2 \rho^2, \quad (6)$$

where $\mathbb{E}(\cdot)$ denotes the expectation,

$$\begin{aligned} \bar{\sigma}^2 &:= \max_{l=1, \dots, L} \lambda_{\max}(\mathbb{E}\{E_l E_l^H\}), \\ \rho^2 &:= \frac{\text{tr}(\sum_{l=1}^L \Psi_l^H \Psi_l)}{\lambda_{\min}^2(\sum_{l=1}^L \Psi_l^H \Psi_l)}, \end{aligned} \quad (7)$$

$\lambda_{\max}(\cdot)$ denotes the largest eigenvalue of its argument, and the expectation is taken over the model mismatch.

Given fixed K and R , it is natural to expect that $\bar{\sigma}^2$ is bounded by some constant independent of L , and so the expected error bound in (6) largely depends on ρ^2 in (7). When training samples are i.i.d., one may further expect $(1/L)\sum_{l=1}^L \Psi_l^H \Psi_l$ to concentrate around its expectation, roughly resulting in $\rho^2 \propto 1/L$, with a proportionality constant that depends on R and the statistics of the training data. Fig. 1 illustrates ρ^2 for various image datasets, providing empirical evidence of this decrease in real data.

Our second theorem provides a *probabilistic* error bound via concentration inequalities, given i.i.d. training sample and model mismatch pairs $(x_1, E_1), \dots, (x_L, E_L)$.¹ It removes the zero-mean assumption for the model mismatches $\{E_l: \forall l\}$ in Corollary 2 that might be strong, e.g., if training data are not preprocessed to have zero mean.

Theorem 3.—Suppose that training sample and model mismatch pairs

$$(x_1, E_1), \dots, (x_L, E_L) \stackrel{iid}{\sim} (x, E), \text{ where } x \text{ and } E \text{ are almost surely bounded, i.e.,}$$

¹We follow the natural convention in sample size analyses of assuming that $\{x_l: \forall l\}$ are i.i.d. samples from an underlying training distribution; see the references cited in Section IV and [5], [6] for other examples. Model mismatches $\{E_l: \forall l\}$ also become i.i.d. across samples at all iterations of CAOL, if “fresh” training samples are used for each update, e.g., as can be done when solving (P1) via mini-batch stochastic optimization.

$$\|x\|_2 \leq \gamma \text{ and } \|E\|_F \leq \sigma, \quad (8)$$

and the matrices in (4) are almost surely full row rank. Then, for any $0 < \delta < \lambda_{\min}(\bar{\Lambda})/(2R\gamma^2)$, the solution D^* to (P1) has error with respect to D_{true} bounded as

$$\|D^* - D_{\text{true}}\|_F^2 \leq 5 \left\{ \frac{\sigma\sqrt{\text{tr}(\bar{\Lambda})/L} + \|\mathbb{E}(\Psi^H E)\|_F + 2\sigma\gamma\sqrt{R}\delta}{\lambda_{\min}(\bar{\Lambda}) - 2\gamma^2 R\delta} \right\}^2, \quad (9)$$

with probability at least

$$1 - 3R \exp\left(-L \frac{\delta^2/2}{3 + \delta/3}\right), \quad (10)$$

where $\bar{\Lambda} := \mathbb{E}(\Psi^H \Psi)$ and Ψ is constructed from x as in (2).

Taking δ sufficiently small, the high probability error bound (9) is primarily driven by

$$\bar{\rho} := \frac{\sqrt{\text{tr}(\bar{\Lambda})/L}}{\lambda_{\min}(\bar{\Lambda})} \text{ and } \bar{\chi} := \frac{\|\mathbb{E}(\Psi^H E)\|_F}{\lambda_{\min}(\bar{\Lambda})}, \quad (11)$$

where $\bar{\rho}$ is analogous to ρ in (7), and $\bar{\chi}$ captures how correlated the model mismatch is to the training samples. As the number L of training samples increases, $\bar{\rho}$ decreases as $1/\sqrt{L}$. On the other hand, $\bar{\chi}$ is constant with respect to L and provides a floor for the bound. Fig. 2 illustrates $\bar{\chi}$ for CAOL iterates from different image datasets, and provides empirical evidence that this term can indeed be small in real data. If the model mismatch is sufficiently uncorrelated with the training samples, i.e., $\bar{\chi}$ is practically zero, then only the $\bar{\rho}$ term remains and this term decreases with L . Namely, if model mismatch is entirely uncorrelated with the training samples, then using many samples decreases the error bound to (effectively) zero.

IV. Related Works

Sample complexity [7] and synthesis (or reconstruction) error [8] have been studied in the context of synthesis operator learning (e.g., dictionary learning [9]); see the cited papers and references therein. A similar understanding for (C)AOL has however remained largely open; existing works focus primarily on establishing (C)AOL models and their algorithmic challenges [1], [10]–[13]. The authors in [14] studied sample complexity for a patch-based AOL method, but the form of their model differs from that of ours (P0). Specifically, they consider the following AOL problem: $\min_D \sum_l f(D^T \hat{x}_l) + g(D)$, where $f(\cdot)$ is a sparsity promoting function (e.g., a smooth approximation of the ℓ_1 -quasi-norm [14]), $g(\cdot)$ is a regularizer or constraint for the filter matrix D , and $\{\hat{x}_l : l = 1, \dots, L\}$ is a set of training *patches* (not images).

V. Proof of Theorem 1

Rewriting (P1) yields that D^* is a solution of the (scaled) orthogonal Procrustes problem [1, §S.VII]:

$$\operatorname{argmin}_D \|\widetilde{\Psi}D - \widetilde{Z}\|_F^2, \text{ subj. to } DD^H = \frac{1}{R} \cdot I, \quad (12)$$

where $\widetilde{\Psi} \in \mathbb{C}^{LN \times R}$ arises by stacking Ψ_1, \dots, Ψ_L vertically and $\widetilde{Z} \in \mathbb{C}^{LN \times K}$ arises likewise from Z_1, \dots, Z_L . Similarly, since $\Psi_l D_{\text{true}} = Z_{\text{true},l}$ as in (3), D_{true} is a solution of the analogous (scaled) orthogonal Procrustes problem

$$\operatorname{argmin}_D \|\widetilde{\Psi}D - \widetilde{Z}_{\text{true}}\|_F^2, \text{ subj. to } DD^H = \frac{1}{R} \cdot I, \quad (13)$$

where $\widetilde{Z}_{\text{true}} \in \mathbb{C}^{LN \times K}$ arises by stacking $Z_{\text{true},1}, \dots, Z_{\text{true},L}$ vertically.

By assumption, both $\widetilde{\Psi}^H \widetilde{Z}$ and $\widetilde{\Psi}^H \widetilde{Z}_{\text{true}}$ are full row rank and so (12) and (13) have unique solutions given by the unique (scaled) polar factors

$$D^* = \frac{1}{\sqrt{R}} Q(\widetilde{Z}^H \widetilde{\Psi})^H D_{\text{true}} = \frac{1}{\sqrt{R}} Q(\widetilde{Z}_{\text{true}}^H \widetilde{\Psi})^H \quad (14)$$

where $Q(\cdot)$ denotes the polar factor of its argument, and can be computed as $Q(A) = WV^H$ from the (thin) singular value decomposition $A = W\Sigma V^H$.

Thus we have

$$\begin{aligned} & \|D^* - D_{\text{true}}\|_F^2 \\ &= \frac{1}{R} \left\| Q(\widetilde{Z}_{\text{true}}^H \widetilde{\Psi}) - Q(\widetilde{Z}^H \widetilde{\Psi}) \right\|_F^2 \\ &\leq \frac{1}{R} \|\widetilde{E}^H \widetilde{\Psi}\|_F^2 \left\{ \left[\frac{2}{\sigma_R(\widetilde{Z}_{\text{true}}^H \widetilde{\Psi}) + \sigma_R(\widetilde{Z}^H \widetilde{\Psi})} \right]^2 + \left[\frac{1}{\max\{\sigma_R(\widetilde{Z}_{\text{true}}^H \widetilde{\Psi}), \sigma_R(\widetilde{Z}^H \widetilde{\Psi})\}} \right]^2 \right\} \\ &\leq \frac{1}{R} \|\widetilde{E}^H \widetilde{\Psi}\|_F^2 \left\{ \left[\frac{2}{\sigma_R(\widetilde{Z}_{\text{true}}^H \widetilde{\Psi})} \right]^2 + \left[\frac{1}{\sigma_R(\widetilde{Z}_{\text{true}}^H \widetilde{\Psi})} \right]^2 \right\} \\ &= \frac{5}{R} \frac{\|\widetilde{\Psi}^H \widetilde{E}\|_F^2}{\sigma_R^2(\widetilde{\Psi}^H \widetilde{Z}_{\text{true}})} = \frac{5}{R} \frac{\|\sum_{l=1}^L \Psi_l^H E_l\|_F^2}{\sigma_R^2(\sum_{l=1}^L \Psi_l^H Z_{\text{true},l})} \end{aligned} \quad (15)$$

where $\widetilde{E} = \widetilde{Z} - \widetilde{Z}_{\text{true}}$ is exactly E_1, \dots, E_L stacked vertically, and $\sigma_r(\cdot)$ denotes the r th largest singular value of its argument. The first inequality holds by the perturbation bound in [15, Thm. 3], and the second holds since $\sigma_R(\widetilde{Z}^H \widetilde{\Psi}) \geq 0$. Recalling that $Z_{\text{true},l} = \Psi_l D_{\text{true}}$, we rewrite the denominator of (15) as

$$\begin{aligned}
\sigma_R^2 \left(\sum_{l=1}^L \Psi_l^H Z_{\text{true}, l} \right) &= \sigma_R^2 \left(\sum_{l=1}^L \Psi_l^H \Psi_l D_{\text{true}} \right) \\
&= \frac{1}{R} \sigma_R^2 \left(\sum_{l=1}^L \Psi_l^H \Psi_l \right) = \frac{1}{R} \lambda_{\min}^2 \left(\sum_{l=1}^L \Psi_l^H \Psi_l \right),
\end{aligned} \tag{16}$$

where the second equality holds because $D_{\text{true}} D_{\text{true}}^H = (1/R)I$. Substituting (16) into (15) yields (5).

VI. Proof of Corollary 2

Taking the expectation of (5) over the model mismatch amounts to taking the expectation of the numerator of the upper bound in (5):

$$\begin{aligned}
&\mathbb{E} \left\| \sum_{l=1}^L \Psi_l^H E_l \right\|_F^2 \\
&= \sum_{l=1}^L \mathbb{E} \|\Psi_l^H E_l\|_F^2 = \sum_{l=1}^L \text{tr}(\Psi_l^H \mathbb{E}\{E_l E_l^H\} \Psi_l) \\
&\leq \sum_{l=1}^L \lambda_{\max}(\mathbb{E}\{E_l E_l^H\}) \cdot \|\Psi_l\|_F^2 \leq \bar{\sigma}^2 \cdot \sum_{l=1}^L \|\Psi_l\|_F^2,
\end{aligned} \tag{17}$$

where the first equality holds by using the assumption that E_l is zero-mean and independent over l , the second equality follows by expanding the Frobenius norm then applying linearity of the trace and expectation, the first inequality holds since $v^H M v \leq \lambda_{\max}(M) \cdot \|v\|_2^2$ for any vector v and Hermitian matrix M , and the last inequality follows from the definition of $\bar{\sigma}^2$. Rewriting (17) using the identity $\sum_{l=1}^L \|\Psi_l\|_F^2 = \sum_{l=1}^L \text{tr}(\Psi_l^H \Psi_l) = \text{tr}(\sum_{l=1}^L \Psi_l^H \Psi_l)$ yields the result (6).

VII. Proof of Theorem 3

We derive two high probability bounds, one each for the numerator and denominator of (5). Then, the bound (9) with probability (10) follows by combining the two via a union bound. Before we begin, note that (8) implies that $\|\Psi\|_2 \leq \|\Psi\|_F \leq \gamma\sqrt{R}$ almost surely; our proofs use this inequality multiple times.

A. Upper bound for numerator

Observe first that

$$\begin{aligned}
&\left\| \sum_{l=1}^L \Psi_l^H E_l \right\|_F = \left\| L \mathbb{E}(\Psi_l^H E_l) + \sum_{l=1}^L \{\Psi_l^H E_l - \mathbb{E}(\Psi_l^H E_l)\} \right\|_F \leq L \|\mathbb{E}(\Psi_l^H E_l)\|_F \\
&+ \left\| \sum_{l=1}^L \xi_l \right\|_2,
\end{aligned} \tag{18}$$

where $\xi_l := \text{vec}\{\Psi_l^H E_l - \mathbb{E}(\Psi_l^H E_l)\} \in \mathbb{C}^{RK}$ for $l = 1, \dots, L$. We next bound $\|\sum_{l=1}^L \xi_l\|_2$ via the vector Bernstein inequality [16, Cor. 8.44]. Note that ξ_1, \dots, ξ_L are i.i.d. with $\mathbb{E}\xi_l = 0$ (by construction). Furthermore, ξ_l is almost surely bounded as

$$\begin{aligned} \|\xi_l\|_2 &= \|\Psi_l^H E_l - \mathbb{E}(\Psi_l^H E_l)\|_F \\ &\leq \|\Psi_l^H E_l\|_F + \|\mathbb{E}(\Psi_l^H E_l)\|_F \text{ (Triangle ineq.)} \\ &\leq \|\Psi_l^H E_l\|_F + \|\mathbb{E}(\Psi_l^H E_l)\|_F \text{ (Jensen's ineq.)} \\ &\leq \|\Psi_l\|_F \|E_l\|_F + \mathbb{E}\|\Psi_l\|_F \|E_l\|_F \\ &\leq 2\sigma\gamma\sqrt{R}. \end{aligned}$$

Thus the vector Bernstein inequality [16, Cor. 8.44] yields that for any $t > 0$,

$$\left\| \sum_{l=1}^L \xi_l \right\|_2 \leq \sigma\sqrt{L}\sqrt{\text{tr}(\bar{\Lambda})} + t, \quad (19)$$

with probability at least

$$1 - \exp\left\{ \frac{-t^2/2}{3L(2\sigma\gamma\sqrt{R})^2 + t(2\sigma\gamma\sqrt{R})/3} \right\}. \quad (20)$$

We obtained (19) by the following simplification:

$$\begin{aligned} \mathbb{E} \left\| \sum_{l=1}^L \xi_l \right\|_2 &\leq \sqrt{\mathbb{E} \left\| \sum_{l=1}^L \xi_l \right\|_2^2} = \sqrt{L\mathbb{E}\|\xi_l\|_2^2} \\ &= \sqrt{L\mathbb{E}\|\Psi_l^H E_l - \mathbb{E}(\Psi_l^H E_l)\|_F^2} \\ &= \sqrt{L\left\{ \mathbb{E}\|\Psi_l^H E_l\|_F^2 - \|\mathbb{E}(\Psi_l^H E_l)\|_F^2 \right\}} \\ &\leq \sqrt{L\mathbb{E}\|\Psi_l^H E_l\|_F^2} \leq \sqrt{L\mathbb{E}(\|\Psi_l\|_F^2 \|E_l\|_F^2)} \\ &\leq \sqrt{L\sigma^2\mathbb{E}\|\Psi_l\|_F^2} = \sigma\sqrt{L}\sqrt{\text{tr}(\bar{\Lambda})}, \end{aligned}$$

where the third equality holds by

$$\mathbb{E}\|A - \mathbb{E}A\|_F^2 = \sum_{i,j} \mathbb{E}(A_{i,j} - \mathbb{E}A_{i,j})^2 = \sum_{i,j} \mathbb{E}A_{i,j}^2 - (\mathbb{E}A_{i,j})^2 = \mathbb{E}\|A\|_F^2 - \|\mathbb{E}A\|_F^2.$$

We obtained (20) by the following simplifications:

$$\sup_{\|x\|_2 \leq 1} \mathbb{E}|x^H \xi_l|^2 \leq \mathbb{E}\|\xi_l\|_2^2 \leq (2\sigma\gamma\sqrt{R})^2,$$

$$\mathbb{E} \left\| \sum_{l=1}^L \xi_l \right\|_2 \leq \mathbb{E} \sum_{l=1}^L \|\xi_l\|_2 \leq L\mathbb{E}\|\xi_l\|_2 \leq L(2\sigma\gamma\sqrt{R}).$$

Applying (19) and (20) with $t = 2\sigma\gamma\sqrt{R}L\delta$ to the square of (18) yields

$$\left\| \sum_{l=1}^L \Psi_l^H E_l \right\|_F^2 \leq L^2 \left\{ \sigma \sqrt{\text{tr}(\bar{\Lambda})/L} + \left\| \mathbb{E}(\Psi_l^H E_l) \right\|_F + 2\sigma\gamma\sqrt{R\delta} \right\}^2, \quad (21)$$

with probability at least $1 - \exp\left(-L \frac{\delta^2/2}{3 + \delta/3}\right)$.

B. Lower bound for denominator

Observe that $\sum_{l=1}^L \Psi_l^H \Psi_l = L\bar{\Lambda} + \sum_{l=1}^L \Lambda_l$, where $\Lambda_l := \Psi_l^H \Psi_l - \bar{\Lambda}$, so Weyl's inequality [17] yields

$$\lambda_{\min}\left(\sum_{l=1}^L \Psi_l^H \Psi_l\right) \geq \lambda_{\min}(L\bar{\Lambda}) - \left\| \sum_{l=1}^L \Lambda_l \right\|_2, \quad (22)$$

and it remains to bound $\left\| \sum_{l=1}^L \Lambda_l \right\|_2$. We do so by using the Matrix Bernstein inequality [16, Cor. 8.15].

Note that $\Lambda_1, \dots, \Lambda_L$ are i.i.d. (since x_1, \dots, x_L are i.i.d.) and $\mathbb{E}\Lambda_l = 0$. Furthermore, Λ_l is almost surely bounded as

$$\begin{aligned} \|\Lambda_l\|_2 &= \left\| \Psi_l^H \Psi_l - \mathbb{E}(\Psi_l^H \Psi_l) \right\|_2 \\ &\leq \left\| \Psi_l^H \Psi_l \right\|_2 + \left\| \mathbb{E}(\Psi_l^H \Psi_l) \right\|_2 \text{ (Triangle ineq.)} \\ &\leq \left\| \Psi_l^H \Psi_l \right\|_2 + \mathbb{E} \left\| \Psi_l^H \Psi_l \right\|_2 \text{ (Jensen's ineq.)} \\ &= \|\Psi_l\|_2^2 + \mathbb{E} \|\Psi_l\|_2^2 \leq 2\gamma^2 R. \end{aligned}$$

Thus, the Matrix Bernstein inequality [16, Cor. 8.15] yields that for any $t > 0$,

$$\mathbb{P}\left\{ \left\| \sum_{l=1}^L \Lambda_l \right\|_2 \geq t \right\} \leq 2R \exp\left\{ \frac{-t^2/2}{L(2\gamma^2 R)^2 + 2\gamma^2 R t/3} \right\}, \quad (23)$$

where we use the following simplification:

$$\left\| \sum_{l=1}^L \mathbb{E}\Lambda_l^2 \right\|_2 = L \|\mathbb{E}\Lambda_l^2\|_2 \leq L \mathbb{E} \|\Lambda_l\|_2^2 \leq L(2\gamma^2 R)^2.$$

Applying (23) with $t = 2\gamma^2 RL\delta$ to the square of (22) yields

$$\lambda_{\min}^2\left(\sum_{l=1}^L \Psi_l^H \Psi_l\right) \geq L^2 \left\{ \lambda_{\min}(\bar{\Lambda}) - 2\gamma^2 R\delta \right\}^2. \quad (24)$$

with probability at least $1 - 2R \exp\left(-L \frac{\delta^2/2}{1 + \delta/3}\right)$.

C. Combined bound

Combining the bounds (21) and (24) via a union bound yields (9) with probability at least

$$1 - \exp\left(-L \frac{\delta^2/2}{3 + \delta/3}\right) - 2R \exp\left(-L \frac{\delta^2/2}{1 + \delta/3}\right), \quad (25)$$

which is greater than or equal to (10).

Acknowledgments

This work is supported in part by the Keck Foundation and NIH grant U01 EB018753. BA is supported by NSERC grant 611675.

REFERENCES

- [1]. Chun IY and Fessler JA, "Convolutional analysis operator learning: Acceleration and convergence," submitted, 1 2018 [Online]. Available: <http://arxiv.org/abs/1802.05584>
- [2]. Chun IY and Fessler JA, "Convolutional dictionary learning: Acceleration and convergence," *IEEE Trans. Image Process*, vol. 27, no. 4, pp. 1697–1712, 4 2018. [PubMed: 28991744]
- [3]. Chun IY and Fessler JA, "Convergent convolutional dictionary learning using adaptive contrast enhancement (CDL-ACE): Application of CDL to image denoising," in *Proc. Sampling Theory and Appl. (SampTA)*, Tallinn, Estonia, 7 2017, pp. 460–464.
- [4]. Chun IY and Fessler JA, "Convolutional analysis operator learning: Application to sparse-view CT," in *Proc. Asilomar Conf. on Signals, Syst., and Comput*, Pacific Grove, CA, 10 2018, pp. 1631–1635.
- [5]. Hastie T, Tibshirani R, and Friedman J, *The elements of statistical learning: Data mining, inference, and prediction*, ser. Springer series in statistics. New York, NY: Springer, 2009.
- [6]. Mohri M, Rostamizadeh A, and Talwalkar A, *Foundations of machine learning*. Cambridge, MA: MIT Press, 2018.
- [7]. Shakeri Z, Sarwate AD, and Bajwa WU, "Sample complexity bounds for dictionary learning from vector- and tensor-valued data," in *Information Theoretic Methods in Data Science*, Rodrigues M and Eldar Y, Eds. Cambridge, UK: Cambridge University Press, 2019, ch. 5.
- [8]. Singh S, Poczos B, and Ma J, "Minimax reconstruction risk of convolutional sparse dictionary learning," in *Proc. Int. Conf. on Artif. Int. and Stat*, ser. *Proc. Mach. Learn. Res.*, vol. 84, Playa Blanca, Lanzarote, Canary Islands, 4 2018, pp. 1327–1336.
- [9]. Aharon M, Elad M, and Bruckstein A, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Process*, vol. 54, no. 11, pp. 4311–4322, 11 2006.
- [10]. Yaghoobi M, Nam S, Gribonval R, and Davies ME, "Constrained overcomplete analysis operator learning for cospase signal modelling," *IEEE Trans. Signal Process*, vol. 61, no. 9, pp. 2341–2355, 3 2013.
- [11]. Hawe S, Kleinsteuber M, and Diepold K, "Analysis operator learning and its application to image reconstruction," *IEEE Trans. Image Process*, vol. 22, no. 6, pp. 2138–2150, 6 2013. [PubMed: 23412611]
- [12]. Cai J-F, Ji H, Shen Z, and Ye G-B, "Data-driven tight frame construction and image denoising," *Appl. Comput. Harmon. Anal*, vol. 37, no. 1, pp. 89–105, 10 2014.
- [13]. Ravishankar S and Bresler Y, "ℓ₀ sparsifying transform learning with efficient optimal updates and convergence guarantees," *IEEE Trans. Sig. Process*, vol. 63, no. 9, pp. 2389–2404, 5 2015.
- [14]. Seibert M, Wörmann J, Gribonval R, and Kleinsteuber M, "Learning co-sparse analysis operators with separable structures," *IEEE Trans. Signal Process*, vol. 64, no. 1, pp. 120–130, 1 2016.
- [15]. Li R-C, "New perturbation bounds for the unitary polar factor," *SIAM J. Matrix Anal. Appl*, vol. 16, no. 1, pp. 327–332, 1 1995.

- [16]. Foucart S and Rauhut H, A mathematical introduction to compressive sensing. New York, NY: Springer, 2013.
- [17]. Weyl H, “Das asymptotische verteilungsgesetz der eigenwerte linearer partieller differentialgleichungen (mit einer anwendung auf die theorie der hohlraumstrahlung),” *Mathematische Annalen*, vol. 71, no. 4, pp. 441–479, 12 1912.

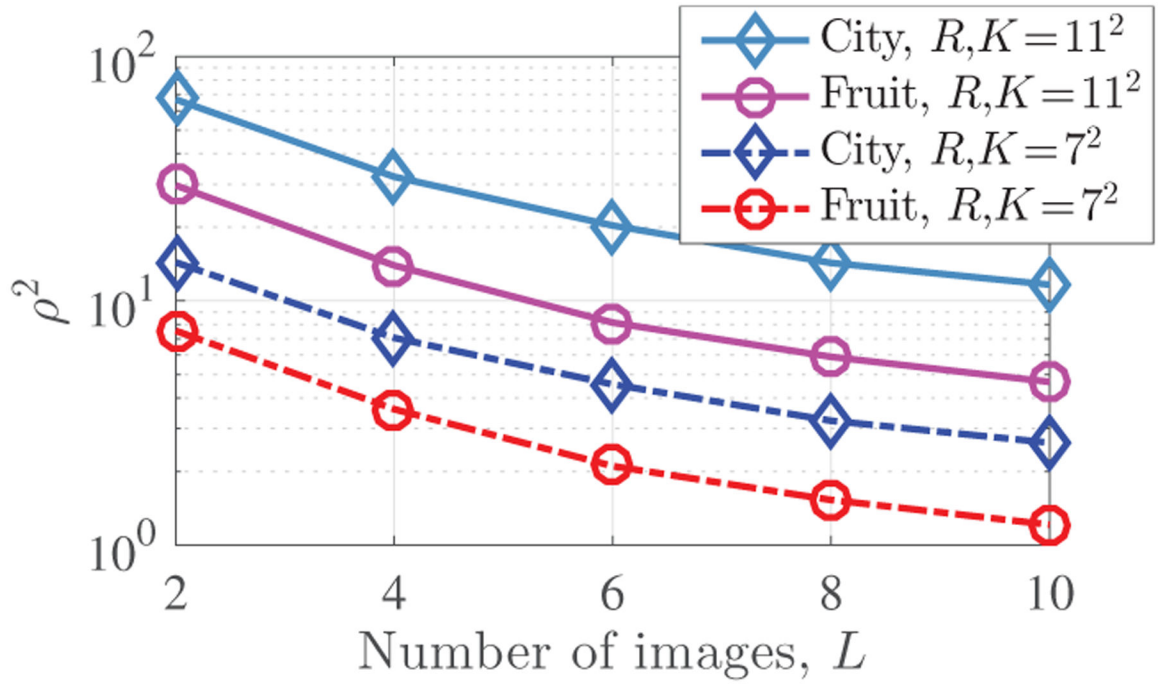


Fig. 1.

Empirical values of ρ^2 in (7) show a decrease with L for different datasets and filter dimensions. (The fruit and city datasets with $L=10$ and $N=10^4$ were preprocessed with contrast enhancement and mean subtraction; see details of datasets and experiments in [1], [2] and references therein. For $L < 10$, the results are averaged over 50 datasets randomly selected from the full datasets.) Under the assumptions of Corollary 2, the decrease in this quantity leads to a better expected error bound in (5). Without preprocessing, the quantity ρ^2 increases by a factor of around 10^3 .

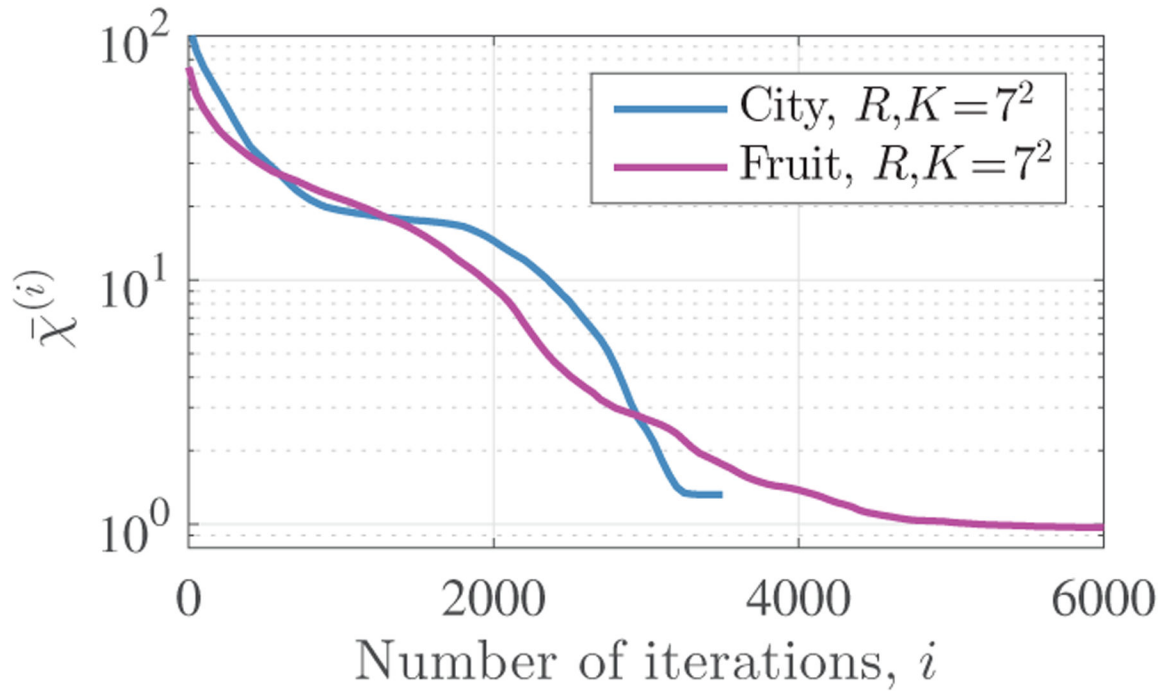


Fig. 2.

Empirical estimate of $\bar{\chi}$ in (11) across iterations in the alternating optimization algorithm [1] that solves CAOL (P0) with $\alpha = 10^{-3}$. (The fruit and city datasets with $L = 10$ and $N = 10^4$ were preprocessed with contrast enhancement and mean subtraction; see details of datasets and experiments in [1], [2] and references therein. The model mismatches $\{E_l^{(i)} : \forall l\}$ at the i th iteration were calculated every 50 iterations based on (3), where we use the converged filters for D_{true} .) Observe that $\bar{\chi}^{(i)}$ generally decreases over iterations; when $\bar{\chi}$ is small, the high probability error bound (9) in Theorem 3 depends primarily on $\bar{\rho}$ defined in (11).