Working Locally Thinking Globally: Guarantees for Convolutional Sparse Coding

Vardan Papyan*, Jeremias Sulam* and Michael Elad
Computer Science Department
Technion - Israel Institute of Technology
vardanp@campus.technion.ac.il, {jsulam,elad}@cs.technion.ac.il

***Abstract*—The acclaimed sparse representation model has led to remarkable results in various signal processing tasks. However, despite its initial purpose of serving as a global prior for entire signals, it has been commonly used for modeling low dimensional patches due to the computational constraints entailed when deployed with learned dictionaries. The emerging convolutional sparse coding (CSC) model comes as a globally-aware alternative. Several works have presented algorithmic solutions to the global pursuit problem under this new model, yet very few truly effective guarantees are known for their success. In this work, we address the theoretical aspects of the CSC model, providing the first meaningful answers to questions of uniqueness of solutions and success of pursuit algorithms. Moreover, we further extend this analysis to the noisy regime, addressing the stability of the sparsest solutions and of the associated algorithms. Finally, we demonstrate practical approaches for solving the global pursuit problem via simple local processing.**

I. Introduction

The sparse representation model assumes a signal $\mathbf{X}$ can be (well) approximated by the product of a dictionary $\mathbf{D}$ and a sparse vector $\mathbf{\Gamma}$. The convolutional counterpart further imposes a specific structure on the matrix $\mathbf{D}$ – that it is composed of shifts of a local dictionary $\mathbf{D}_L$. Recently, this model has been shown to provide a variety of useful applications [1]–[7]. However, the theoretical aspects of this model were disregarded, with the assumption that the original sparse theory holds for this model as well [8]–[10]. These results rely on properties of the dictionary $\mathbf{D}$, such as its mutual coherence, and on the maximal number of total non-zero entries in the global representation vector.

Consider a sparse vector $\mathbf{\Gamma}$ which represents a global (convolutional) signal. Assume further that this vector has a few non-zeros. If these were to be clustered together in a given portion (*stripe*) $\boldsymbol{\gamma}_i$, as the one in Figure 1, the local patch corresponding to this stripe would be very complex, and pursuit methods would likely fail in recovering it. On the contrary, consider the case where these non-zeros are spread all throughout the vector $\mathbf{\Gamma}$. This would clearly imply much simpler local patches, facilitating their successful recovery. This simple example comes to show the futility of the traditional global ℓ_0-norm in the convolutional setting, and it will be the pillar of our intuition throughout our work.

II. From global sparsity to local constraints

The following is a measure that provides a local notion of sparsity within a global sparse vector.

Definition 1. Let the $\ell_{0,\infty}$ pseudo-norm of a global vector $\mathbf{\Gamma}$ be $\|\mathbf{\Gamma}\|_{0,\infty} = \max_i \|\boldsymbol{\gamma}_i\|_0$.

In words, this quantifies the number of non-zeros in the densest stripe γ_i of the global $\mathbf{\Gamma}$. Intuitively, by constraining the $\ell_{0,\infty}$ norm to be low, we are essentially limiting the sparsity of all the stripes $\boldsymbol{\gamma}_i$. Armed with this definition, we define the $P_{0,\infty}$ problem:

$$(P_{0,\infty}): \quad \min_{\mathbf{\Gamma}} \quad \|\mathbf{\Gamma}\|_{0,\infty} \ \text{ s.t. } \ \mathbf{D\Gamma} = \mathbf{X}.$$

* The authors contributed equally to this work.

When dealing with a global signal, instead of solving the P_0 problem as is commonly done, we aim to solve the above defined objective instead. Note that we are not limiting the overall number of zeros in $\mathbf{\Gamma}$, but rather putting a restriction on its local density.

III. Uniqueness and recovery guarantees

Does a unique solution to the $P_{0,\infty}$ problem exist? and under which circumstances? In the first part of our work [11] we show that under simple local sparsity constraints (enforced via the $\ell_{0,\infty}$ norm) the answer to this question is positive. Interestingly, this unique representation can be also be recovered using popular pursuit algorithms such as the Basis Pursuit (BP) and the Orthogonal Matching Pursuit (OMP). This is summarized in the following theorem:

Theorem 2. *Given the system of linear equations $\mathbf{X} = \mathbf{D\Gamma}$, where $\mathbf{D}$ is a convolutional dictionary with mutual coherence $\mu(\mathbf{D})$, if a solution $\mathbf{\Gamma}$ exists satisfying $\|\mathbf{\Gamma}\|_{0,\infty} < \frac{1}{2}\left(1 + \frac{1}{\mu(\mathbf{D})}\right)$, then OMP and BP are guaranteed to recover it.*

Importantly, the local constraint provides global guarantees that scale with the signal dimension, thus yielding useful bounds.

IV. Stability results

In the second part of our work [12] we consider signal perturbations and model deviations by extending the $P_{0,\infty}$ problem to a relaxed $P_{0,\infty}^{\epsilon}$ version. We address questions of stability of the sparsest solutions to this problem and the ability of pursuit algorithms – both greedy and convex – to approximate them. To this end, we generalize classical definitions, such as the RIP, to the convolutional model, and connect existing notions, such as the ERC, to our setting. As expected, from our analysis it follows that meaningful pursuit stability bounds arise only once one assumes locally-bounded noise. Intuitively, if a signal is severely contaminated in a small local area, any global pursuit will likely fail.

V. Practical aspects

Note that the above theoretical results show that global pursuit algorithms are guaranteed to recover the exact (or approximate) solutions to the $P_{0,\infty}$ (or $P_{0,\infty}^{\epsilon}$) problem. However, these problem are often too big to be tackled directly. On the algorithmic side, we demonstrate how to harness the local-global connections in the CSC model to solve global pursuit using only local operations. This offers a first of its kind bridge between global modeling of signals and their patch-based local treatment.

Acknowledgement

The research leading to these results has received funding from the European Research Council under European Unions Seventh Framework Programme, ERC Grant agreement no. 320649

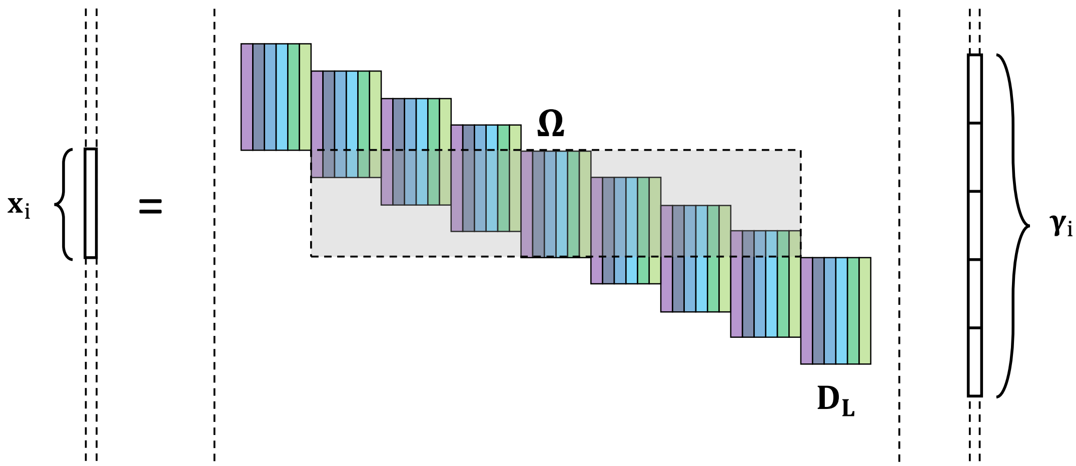

Fig. 1: The i^{th} local system of equations $\mathbf{x}_i = \boldsymbol{\Omega}\boldsymbol{\gamma}_i$, where $\mathbf{x}_i$ is a single patch, $\boldsymbol{\Omega}$ is the stripe dictionary and $\boldsymbol{\gamma}_i$ is a stripe vector.

References

[1] S. Gu, W. Zuo, Q. Xie, D. Meng, X. Feng, and L. Zhang, "Convolutional sparse coding for image super-resolution," in *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1823–1831, 2015.

[2] F. Heide, W. Heidrich, and G. Wetzstein, "Fast and flexible convolutional sparse coding," in *Computer Vision and Pattern Recognition (CVPR), 2015 IEEE Conference on*, pp. 5135–5143, IEEE, 2015.

[3] M. D. Zeiler, D. Krishnan, G. W. Taylor, and R. Fergus, "Deconvolutional networks," in *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pp. 2528–2535, IEEE, 2010.

[4] H. Bristow, A. Eriksson, and S. Lucey, "Fast Convolutional Sparse Coding," *2013 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 391–398, June 2013.

[5] Y. Zhu and S. Lucey, "Convolutional sparse coding for trajectory reconstruction," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 37, no. 3, pp. 529–540, 2015.

[6] B. Wohlberg, "Efficient convolutional sparse coding," in *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*, pp. 7173–7177, IEEE, 2014.

[7] B. Kong and C. C. Fowlkes, "Fast convolutional sparse coding (fcsc)," *Department of Computer Science, University of California, Irvine, Tech. Rep*, 2014.

[8] M. Elad, *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*. Springer Publishing Company, Incorporated, 1st ed., 2010.

[9] J. A. Tropp, "Just Relax : Convex Programming Methods for Identifying Sparse Signals in Noise," *IEEE Transactions on In*, vol. 52, no. 3, pp. 1030–1051, 2006.

[10] J. Tropp, "Greed is Good: Algorithmic Results for Sparse Approximation," *IEEE Transactions on Information Theory*, vol. 50, no. 10, pp. 2231–2242, 2004.

[11] V. Papyan, J. Sulam, and M. Elad, "Working locally thinking globally - part i: Theoretical guarantees for convolutional sparse coding," *CoRR*, vol. abs/1607.02005, 2016.

[12] V. Papyan, J. Sulam, and M. Elad, "Working locally thinking globally - part ii: Stability and algorithms for convolutional sparse coding," *CoRR*, vol. abs/1607.02009, 2016.