Dictionary Recovery through Deep Learning: Simulation & Evaluation Methods

Supervision:

- Thomas Moreau, Parietal, Inria (thomas.moreau@inria.fr)
- Alexandre Gramfort, Parietal, Inria (alexandre.gramfort@inria.fr)

Context Dictionary learning and its convolutional counter part are very popular methods to learn unsupervised and structured representations of a datasets. The core idea is to represent the data points as linear combinations of a few atoms selected from an over-complete family of atoms – also called the dictionary. This dictionary is learned solving an optimization problem that promotes the sparsest representations for the data. This type of method has been applied successfully in many fields, from astrophysics to bio-imaging. In the context of neuroscience, it has been used to develop brain atlases from fMRI [9] or to highlight recurring patterns in MEG [5].

While the problem of learning a dictionary is non-convex, many methods have been proposed to compute an approximate solution. Classical methods typically rely on alternate minimization and online updates [8]. A critical element in these methods is the need to compute a good solution of the sparse coding problem using iterative algorithm. Indeed, this step can require many iterations to output a good enough solution which is needed in order to have a correct gradient. These algorithms have been studied in the literature and it has been shown that under certain assumptions on the generating model for the observation, the dictionary could be retrieved up to permutations [3, 4, 7].

Recent studies have shown that deep neural networks can be designed based on iterative thresholding algorithms used to solve sparse coding. Their use speeds up and improves inference compared to original solvers [6]. Returning to the point of view of an iterative algorithm, it is possible to study these models with the classical tools of convex optimization and sparse coding. Our recent works [1, 2] take this point of view to show some properties of networks based on iterative thresholding. These networks can also be used in order to make the sparse coding step in itself differentiable and facilitate the dictionary learning [10]. While the first empirical approaches have shown the potential of these methods – particularly for denoising tasks – the results are regularly confronted with the problem of overfitting and the evaluation of their quality remains an open question.

Quantifying how well a dictionary is recovered is complicated as it is rare that we have a ground truth to evaluate it. To guarantee that the learned atoms are not artifacts, it is thus necessary to have theoretical evidence that the learning process will not provide spurious atoms. While some guarantees have been provided for classical methods, the problem becomes more complicated when using neural dictionary learning methods, due to approximation in the optimization process. However, leveraging classical results, it would be interesting to evaluate the recovery properties of neural algorithms on the class of dictionary that can theoretically be recovered with classical methods. This would provide a set of hypothesis on the potential theoretical properties that could be proven for these models.

Methods In order to properly answer these questions, we propose to rely on an empirical approach, simulating controlled dictionary recovery problems that match the theoretical assumption for dictionary recovery. The classical and neural algorithms would then be evaluated on their capacity to successfully estimate the simulated dictionary. The generated dictionary will have controlled mutual incoherence or RIP properties. We will aim at developing a proper library to simulate such problems in order to provide a reference framework for evaluating dictionary recovery methods.

Environment The internship will take place in Inria Saclay, in the Parietal team. This is a large team working focused on mathematical methods for statistical modeling of brain function using neuroimaging data (fMRI, MEG, EEG). Particular topics of interest include machine learning techniques, numerical and parallel optimization, applications to human cognitive neuroscience, and scientific software development.

Requirements

- Strong mathematical background. Knowledge in numerical optimization is a plus.
- Good programming skills in Python. Knowledge of a deep-learning library is a plus.

References

[1] Pierre Ablin, Thomas Moreau, Mathurin Massias, and Alexandre Gramfort. Learning step sizes for unfolded sparse coding. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 13100–13110, Vancouver, BC, Canada, 2019.

[2] Pierre Ablin, Gabriel Peyré, and Thomas Moreau. Super-efficiency of automatic differentiation for functions defined as a minimum. In *International Conference on Machine Learning (ICML)*, July 2020.

[3] Alekh Agarwal, Animashree Anandkumar, and Prateek Jain. Learning sparsely used overcomplete dictionaries via alternating minimization. *Journal of Machine Learning Research (JMLR)*, 35:1–15, 2014.

[4] Sanjeev Arora, R Ge, and A Moitra. New Algorithms for Learning Incoherent and Overcomplete Dictionaries. *Journal of Machine Learning Research (JMLR)*, 35:1–20, 2013.

[5] Tom Dupré la Tour, Thomas Moreau, Mainak Jas, and Alexandre Gramfort. Multivariate Convolutional Sparse Coding for Electromagnetic Brain Signals. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3296–3306, Montreal, Canada, 2018.

[6] Karol Gregor and Yann Le Cun. Learning Fast Approximations of Sparse Coding. In *International Conference on Machine Learning (ICML)*, pages 399–406, 2010.

[7] Rémi Gribonval, Rodolphe Jenatton, and Francis R. Bach. Sparse and spurious: Dictionary learning with noise and outliers. *IEEE Transactions on Information Theory*, 61(11):6298–6319, 2015.

[8] Julien Mairal, Francis R. Bach, Jean Ponce, and Guillermo Sapiro. Online dictionary learning for sparse coding. *International Conference on Machine Learning (ICML)*, (September):689–696, 2009.

[9] Arthur Mensch, Julien Mairal, Bertrand Thirion, and Gael Varoquaux. Stochastic Subsampling for Factorizing Huge Matrices. *IEEE Transactions on Signal Processing*, 66(1):113–128, January 2018.

[10] Bahareh Tolooshams, Sourav Dey, and Demba Ba. Scalable convolutional dictionary learning with constrained recurrent sparse auto-encoders. In *IEEE International Workshop on Machine Learning for Signal Processing (MLSP)*, 2018.