Bentov - Benchmarking Total Variation Solvers

Supervision:

- Thomas Moreau, Parietal, Inria (thomas.moreau@inria.fr)
- Alexandre Gramfort, Parietal, Inria (alexandre.gramfort@inria.fr)
- Jean-Christophe Pesquet, CVN, CentraleSupélec (Jean-Christophe.pesquet@centralesupelec.fr)

Context When solving statistical tasks on structured data such as signals or images, it is often necessary to carefully regularize the considered model to boost its statistical power and reduce the effect of the noise. A common regularization in this context is the Total-Variation (TV) regularization, which promotes the sparsity of the signal derivatives, leading to piecewise constant models. For images, this regularisation is typically used for deblurring tasks, as it preserves the objects boundaries [4]. Higher order versions of TV regularization have also been used to capture some more complex regularity, in particular in the context of change point detections or trend filtering [5]. For such regularized models, a large part of the computational cost of statistical learning comes from solving large-scale optimization problems. In order to efficiently solve the problems, many solvers have been specifically developed to allow these methods to scale to very large datasets ([1, 3, 2]). In practice, the choice of one algorithm rather than another depends on many factors, partly related to the data that form the problem and the specific parameters used. However, there are few robust and comprehensive comparisons of these algorithms, with various datasets and parameters. The benchopt library (`https://benchopt.github.io/`), so far mainly developed within the Parietal team, provides a framework to perform these comparisons in a simple, reproducible and shareable way. Benchopt is an open source benchmarking toolbox developed collaboratively, which aims at making easy and reproducible comparison of optimization algorithms (`https://benchopt.github.io/results/`).

Assignment The objective of the internship is to develop a fair and complete framework to compare algorithms for solving problems regularized by total variation for both 1D and 2D problems. As a first step to familiarize with the benchopt library, the trainee will have to create a benchmark for 1D linear regression with TV regularization. He/She will need to define the cost function compatible with the API of the library, propose simulated data of constant time signals by piece and integrate classical resolution methods to the analysis (e.g. proximal gradient descent, ADMM) and to the synthesis (e.g. proximal gradient descent, coordinate descent). This comparison will then be enriched on real data extracted from free databases such as physionet (`https://physionet.org/`). Along the way, the candidate will develop new benchmarking tools for convex optimization algorithms, assessing not only the convergence speed but the quality of the solution in terms of metrics such as sparsity of the solution or the distance to the constraint sets for ADMM solutions. This will ensure that other researchers can easily evaluate the proposed approach and potential subsequent improvements.

Then, he/she will create a benchmark for the more complex problem of 2D TV-regularized linear regression for images. He/She will build on the previous benchmark and integrate a variety of methods, in particular some of the approaches described in [1, 3]. For large images, it is often necessary to rely on specialized hardware such as GPU in order to ensure efficient operations. Algorithms that can be adapted for such computational resources will also be benchmarked on both architectures, in order to assess the performance gain obtained. This will require integration of extra monitoring tools in the benchopt library in order to fairly compare these implementations. Finally, the candidate will also be able to explore resolution methods for higher-order total variation regularizations (which are close variants of the methods for order 1) if time permits.

These developments will ensure that the candidate builds a solid experience on optimization, numerical solvers, and their implementation. Moreover, the candidate will also gain experience on the processes to contribute to an open source package (benchopt), in particular with the use of test driven development, continuous integration, code reviews and version control tools such as git and GitHub.

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] Antonin Chambolle and Thomas Pock. A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, May 2011.

[2] Hamza Cherkaoui, Jeremias Sulam, and Thomas Moreau. Learning to solve TV regularised problems with unrolled algorithms. In *Advances in Neural Information Processing Systems (NeurIPS)*, online, 2020.

[3] Patrick L. Combettes and Jean-Christophe Pesquet. Fixed Point Strategies in Data Science. *IEEE Transactions on Signal Processing*, 69:3878–3905, 2021.

[4] Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1-4):259–268, November 1992.

[5] Ryan J. Tibshirani. Adaptive piecewise polynomial estimation via trend filtering. *The Annals of Statistics*, 42(1), February 2014.