

Mini-Project

(Submission Deadline: May 31st, 2019)

Introduction

The optional mini-project focusses on the practical implementation: We encourage you to investigate the real-world performance of an optimization algorithm for realistic machine learning applications, and thereby gain more insight to that optimization method. The goal of your study is to provide solid empirical evidence for some behaviour aspects (pros and cons) of your chosen optimization algorithm (studied in class or not) on a real machine-learning task. The optimization algorithms can be anything of your choice. Don't be scared to try variants not seen in class, as again the focus of the project is not on theory. Also, any application of relevance to ML is a fine choice (no matter if deep learning, linear models or e.g. random forest).

In other words, your report should look like **the “Experiments”-Section of an ICML, NIPS or ICLR machine learning paper**. You are not asked to invent new algorithms or convergence theory, but only to extend knowledge and insights about their experimental performance, and possibly interpret this in view of existing theory.

Grading. The project is optional and done in groups of 2-3 students. If you decide to do the project in a group, and if your project grade exceeds your exam grade, it will count 20% to your final grade. Project reports (3 page PDF) and code are due May 31st.

Topic Examples

The list below is by no means exhaustive, and you are encouraged to freely chose very different aspects and topics as long as they concern sth about optimization, and as long as you study them for any ML relevant application.

- Local minima for deep learning: Can you find differences between the ‘shape’ of local minima that SGD finds, depending on different step-sizes or mini-batch size, vs e.g. AdaGrad or full gradient descent?
- Same thing for GANs? Or for matrix factorizations?
- Along a training trajectory for a deep net, does the behaviour of each step resemble the convex case, is it different early or late in training?
- How do different optimization variants affect generalization (test error)?
- Second-order methods: Do (Quasi-)Newton methods go to differently shaped local minima in neural networks? Or: Is the secant method a viable alternative training method?
- Quantized SGD: As quantization of gradients or iterates becomes more coarse, what happens to the optimization algorithms, and to generalization? Is it different for DL as compared to linear ML models?
- AdaGrad / Adam / signSGD: Can you suggest/try different data-dependent coordinate-wise learning rate schemes and compare them?
- For training deep nets with very large SGD mini-batches: when does the scalability gain disappear? Is it influenced by other properties of the optimizer? For example, what is the effect of first slowly growing the learning rate and then later decreasing it?
- Meta-Learning: Can you learn the learning rate? The direction or curvature? See also the larger field of AutoML.

- How well do zero-order optimization methods do for ML applications, compared to standard first-order methods?
- Distributed or Decentralized SGD: How does the communication topology (and number of machines) influence practical convergence? (experiments can be simulated)
- Asynchronous SGD: How do different delays affect convergence? How does it interplay with momentum? Does it act as a regularizer, like drop-out?
- Some of the theoretical behaviors we analyzed in theory: Do they actually appear in practice or not, for convex models or deep nets?
- Frank-Wolfe or Coordinate-Descent for deep learning?
- ...

Deliverables

- **Written Report.** You will write a maximum 3 page PDF report on your findings, using LaTeX. You can use unlimited additional pages for references and for an appendix. The main paper should be self-contained!
- **Code.** In a language of your choice. Python with PyTorch is recommended for convenient access to gradients and optimizers. External libraries and existing implementations are allowed, if properly cited. You submit the complete executable and documented code, as one .zip file. Rules for the code part:
 - *Reproducibility:* In your submission, provide a script e.g. `run.py` or notebook which produces all results and plots you use in the paper.
 - *Documentation:* Your system must be clearly described in your PDF report and also well-documented in the code itself. A useful ReadMe file must be provided.

Submission URL: <http://epfm117.hotcrp.com>

Submission deadline: May 31st, 2019 (midnight)

Grading Criteria

We will grade you on the scientific contribution you made, that is on the insights gained compared to standard baseline methods. This is only possible based on a rigorous and fair empirical comparison. The criteria are

- **Solid comparison baselines supporting your claims**

Quantify the benefits of your method by providing clear quality measurements of the most important aspects and additions you chose for your model. Start with a very basic baseline, and demonstrate what improvements your contributions yield.

- **Reproducibility**

Your readers should be able to reproduce your results based on your report only. Describe what hyper-parameters you selected and why, chose realistic and representative datasets, and clearly describe the overall pipeline you used.

- **Scientific novelty and creativity**

You will likely be using more than the standard methods we saw in the course. Make sure that your report addresses the following points.

- What is the *specific* aspect which you study, and why this is interesting and important.
- Do a search for related work. Have similar experiments appeared in the literature? If yes, how are your experiments adding additional insight? Discuss the pros/cons of the existing studies compared to your approach.
- How is the algorithm variant or the aspect of your choice helping for optimization speed, accuracy or generalization? E.g. compare the optimization error before and after for the elements of your study.

- Writeup quality

Some advice when writing a scientific report:

- Try to convey a clear story giving the most relevant aspects of your findings. Learning what has not worked can additionally help the reader (and help them better understand *why* you have made the many choices you did), but focus on what is most relevant and interesting.
- Before the submission, have an external person proofread your report. Use a spell-checker.
- Plots are great way to share information that might be hard to convey by writing. Make sure that your plots are understandable, have labels for axes, correct axes limits, and a description.

As usual, your code and report will be automatically checked for plagiarism.

Report Guidelines

Clearly describe your used methods, state your conclusions and argue that the results you obtained make (or do not make) sense, and the reasons behind it. Keep the report short and to the point, with a strict limit of 3 pages. References are allowed to be put on unlimited extra pages, as well as if you need to put additional technical details or writings.

To get started more easily with writing the report, we provide a LaTeX template here

github.com/epfml/OptML_course/tree/master/labs/mini-project/latex-example-paper

The file also contains some more helpful information on how to write a scientific report or paper. We will also help you during the exercise session and office hours if you ask us.

For more guidelines on what makes a good report, see the grading criteria above. In particular, don't forget to take care about

- *Reproducibility*: Not only in the code, but also in the report, do include complete details about each algorithm you tried, e.g. what lambda values you used for ridge regression?
- *Baselines*: Give clear experimental evidence: It is crucial to always report relative differences in the evaluation metrics, that is with and without the element you study, and to include several properly implemented baseline algorithms as a comparison to your approach.

Some additional resources on LaTeX:

- <https://github.com/VoLuong/Begin-Latex-in-minutes> - getting started with LaTeX
- <http://www.maths.tcd.ie/~dwilkins/LaTeXPrimer/> - tutorial on LaTeX
- <http://www.stdout.org/~winston/latex/latexsheet-a4.pdf> - cheat sheet collecting most of all useful commands in LaTeX
- <http://mirror.switch.ch/ftp/mirror/tex/info/first-latex-doc/first-latex-doc.pdf> - example how to create a document with Latex
- <http://en.wikibooks.org/wiki/LaTeX> - detailed tutorial on LaTeX

Producing figures for LaTeX in Python

There are some good visualization tools in Python. "matplotlib" is probably the single most used Python package for 2D-graphics. The relevant tutorials are as follow:

- Matplotlib tutorial: <http://www.labri.fr/perso/nrougier/teaching/matplotlib/>
- Matplotlib tutorial: <https://sites.google.com/site/scigraphs/tutorial>
- Matplotlib Tutorial: http://jakevdp.github.io/mpl_tutorial/

Regarding other useful Python data visualization libraries, please refer to this blog for more information.